AMI Configuration Commands Reference Manual

Part 1

MANU0265-04 - Rev. A - 8/17/98

Software Version 5.3.x

FORE Systems, Inc.

1000 FORE Drive
Warrendale, PA 15086-7502
Phone: 724-742-4444
FAX: 724-742-7742

http://www.fore.com

Legal Notices

FCC CLASS A NOTICE

WARNING: Changes or modifications to this unit not expressly approved by the party responsible for compliance could void this user's authority to operate this equipment.

NOTE: The TNX-210, TNX-1100, ASX-200WG, ASX-200BX, ASX-1000, ASX-4000, *ForeRunnerLE* 25, and *ForeRunnerLE* 155 have been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of the equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

DOC CLASS A NOTICE

This digital apparatus does not exceed Class A limits for radio noise emission for a digital device as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le present appareil numerique n'emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de la class A prescrites dans le reglement sur le brouillage radioelectrique edicte par le ministere des Communications du Canada.

VCCI CLASS 1 NOTICE

　この装置は、第一種情報処理装置（商工業地域において使用されるべき情報処理装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会(VCCI)基準に適合しております。
　従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。
　取扱説明書に従って正しい取り扱いをして下さい。

This equipment is in the Class 1 category (Information Technology Equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council For Interference by Information Technology Equipment aimed at preventing radio interference in commercial and/or industrial areas.Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, etc. Read the instructions for correct handling.

FCC REQUIREMENTS (Notice to Users of DS1 Service)

The following instructions are provided to ensure compliance with the Federal Communications Commission (FCC) Rules, Part 68.

(1) This device must only be connected to the DS1 network connected behind an FCC Part 68 registered channel service unit. Direct connection is not allowed.

(2) Before connecting your unit, you must inform the telephone company of the following information:

Port ID	REN/SOC	FIC	USOC
NM-6/DS1C NM-2/DS1C NM-8/DS1D NM-4/DS1D	6.0N	04DU9-BN, 04DU9-DN, 04DU9-1ZN, and 04DU9-1SN	RJ48C

(3) If the unit appears to be malfunctioning, it should be disconnected from the telephone lines until you learn if your equipment or the telephone line is the source of the trouble. If your equipment needs repair, it should not be reconnected until it is repaired.

(4) If the telephone company finds that this equipment is exceeding tolerable parameters, the telephone company can temporarily disconnect service, although they will attempt to give you advance notice if possible.

(5) Under the FCC Rules, no customer is authorized to repair this equipment. This restriction applies regardless of whether the equipment is in or out of warranty.

(6) If the telephone company alters their equipment in a manner that will affect use of this device, they must give you advance warning so as to give you the opportunity for uninterrupted service. You will be advised of your right to file a complaint with the FCC.

CANADIAN IC CS-03 COMPLIANCE STATEMENT

NOTICE: The Industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational and safety requirements. The Industry Canada label does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. In some cases, the company's inside wiring associated with a single line individual service may be extended by means of a certified connector assembly (telephone extension cord). The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs to certified equipment should be made by an authorized Canadian maintenance facility designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important in rural areas.

Caution: Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

E1 AND E3 NOTICE

The E1 (NM-6/E1C, NM-2/E1C, NM-8/E1D, and NM-4/E1D) and E3 (NM-4/E3C, NM-2/E3C, NM-4/E3D, and NM-2/E3D) network modules that are described in this manual are approved for use in FORE Systems' host systems providing that the instructions below are strictly observed. Failure to follow these instructions invalidates the approval.

Pan European Approval - CE Marking

Pan European approval of the E1 network module was issued by BABT following assessment against CTR12. This means that it can be connected to ONP and unstructured PTO-provided private circuits with 120 Ω interfaces in all European countries, according to Telecommunications Terminal Equipment (TTE) Directive 91/263/EEC. Thus, the following CE mark applies:

CE168X

The E1 and E3 network modules conform to safety standard EN60950: 1992 following the provisions of Low Voltage Product Safety Directive 73/23/EEC and CE Marking Directive 93/68/EEC, and can be marked accordingly with the CE symbol.

The E1 and E3 network modules conform to EN55022: 1994 and EN50082-1: 1992 following the provisions of the EMC Directive 89/336/EEC, and can be marked accordingly with the CE symbol.

National Approvals

UK

Network Module	Connects to	Approval Number
E1	PTO-provided private circuits with 75 Ω interfaces	AA60953
E3	PTO-provided private circuits with 75 Ω interfaces	NS/4387/1/T/605954
CEM E1	PTO-provided private circuits with 75 Ω or 120 Ω unstructured interfaces	AA607478

Required User Guide Statements - UK Installation

The network modules are designed for use only with FORE Systems ATM Switches. Use of the network modules in any product not listed in this manual may result in a hazard and will invalidate the regulatory approval. The network modules must be installed in accordance with the installation instructions provided.

The following table shows the available ports and their safety status:

Ports	Safety Status
E1 and E3 Ports	TNV operating at SELV
Bus Connector	SELV

CE NOTICE

Marking by the symbol **CE** indicates compliance of this system to the EMC (Electromagnetic Compatibility) directive of the European Community and compliance to the Low Voltage (Safety) Directive. Such marking is indicative that this system meets or exceeds the following technical standards:

- EN 55022 - "Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment."
- EN 50082-1 - "Electromagnetic compatibility - Generic immunity standard Part 1: Residential, commercial, and light industry."
- IEC 1000-4-2 - "Electromagnetic compatibility for industrial-process measurement and control equipment Part 2: Electrostatic discharge requirements."
- IEC 1000-4-3 - "Electromagnetic compatibility for industrial-process measurement and control equipment Part 3: Radiate electromagnetic field requirements."
- IEC 1000-4-4 - "Electromagnetic compatibility for industrial-process measurement and control equipment Part 4: Electrical fast transient/burst requirements."

SAFETY CERTIFICATIONS

ETL certified to meet Information Technology Equipment safety standards UL 1950, CSA 22.2 No. 950, and EN 60950.

Table of Contents

CHAPTER 4   ATMR Configuration Commands

Table of Contents

CHAPTER 15   NSAP Configuration Commands

Preface

This manual provides the technical information needed to configure the ATM Management Interface (AMI) for the *ForeRunner*^{TM} family of ATM Switches, TNX ATM Switches, and *ForeRunnerLE* Switches. This manual contains a text and graphical description of the configuration level AMI commands and menus. It describes the commands and menus from `configuration alarms>` to `configuration nsap>`. For a description of the configuration level AMI commands and menus from `configuration port>` to `configuration vpt>`, please see Part 2 of this manual. This document was created for users with various levels of experience. If you have any questions or problems, please contact the FORE Systems' Technical Assistance Center (TAC).

Chapter Summaries

Chapter 1 - Introduction - Provides an overview of the AMI commands and menus that are described in this manual.

Chapter 2 - Alarms Configuration Commands - Describes the `configuration alarms` level AMI commands and menus.

Chapter 3 - ATMARP Configuration Commands - Describes the `configuration atmarp` level AMI commands and menus.

Chapter 4 - ATMR Configuration Commands - Describes the `configuration atmroute` level AMI commands and menus.

Chapter 5 - Board Configuration Commands - Describes the `configuration board` level AMI commands and menus.

Chapter 6 - CEC Configuration Commands - Describes the `configuration cec` level AMI commands and menus.

Chapter 7 - CES Configuration Commands - Describes the `configuration ces` level AMI commands and menus.

Chapter 8 - Frame Relay Commands - Describes the `configuration fratm` level AMI commands and menus.

Chapter 9 - FUNI Configuration Commands - Describes the `configuration funi` level AMI commands and menus.

Chapter 10 - HTTP Configuration Commands - Describes the `configuration http` level AMI commands and menus.

Chapter 11 - ILMI SNMP Proxy Commands - Describes the `configuration ilmiproxy` level AMI commands and menus.

Chapter 12 - IP Configuration Commands - Describes the `configuration ip` level AMI commands and menus.

Chapter 13 - LANE Configuration Commands - Describes the `configuration lane` level AMI commands and menus.

Chapter 14 - Network Module Commands - Describes the `configuration module` level AMI commands and menus.

Chapter 15 - NSAP Configuration Commands - Describes the `configuration nsap` level AMI commands and menus.

Related Manuals

References are made in this manual to the following manuals:

AMI Configuration Command Reference Manual, Part 2 - Describes the configuration level AMI commands and menus from `configuration port>` to `configuration vpt>`.

ATM Management Interface (AMI) Manual - Describes the root, display, operation, and statistics level AMI commands and menus.

ATM Switch Diagnostics and Troubleshooting Manual - Describes the debug level AMI commands and menus. Also, describes error messages, loopbacks, SCP diagnostics, and ATM Forum PNNI debugging information.

ATM Switch Network Configuration Manual - Discusses topics such as LAN Emulation, Classical IP, ATM Forum PNNI, and *ForeThought* PNNI.

These manuals can be found on the CD and can be read and printed using Acrobat Reader which is also included on the CD. If Acrobat Reader is installed locally, run Acrobat and open the manual from the /DOCS directory of the CD. If Acrobat Reader is not installed locally, run the Acrobat installer to load Acrobat Reader on your machine. Then run the ACROREAD.EXE file in the /DOCS directory of the CD.

Technical Support

In the U.S.A., customers can reach FORE Systems' Technical Assistance Center (TAC) using any one of the following methods:

1. Select the "Support" link from FORE's World Wide Web page:

    http://www.fore.com/

2. Send questions, via e-mail, to:

    support@fore.com

3. Telephone questions to "support" at:

    800-671-FORE (3673) or **724-742-6999**

4. FAX questions to "support" at:

    724-742-7900

Technical support for customers outside the United States should be handled through the local distributor or via telephone at the following number:

+1 724-742-6999

No matter which method is used to reach FORE Support, customers should be ready to provide the following:

*   A support contract ID number
*   The serial number of each product in question
*   All relevant information describing the problem or question

Typographical Styles

Throughout this manual, all specific commands meant to be entered by the user appear on a separate line in bold typeface. In addition, use of the Enter or Return key is represented as <ENTER>. The following example demonstrates this convention:

cd /usr <ENTER>

File names that appear within the text of this manual are represented in the following style: "...the fore_install program installs this distribution."

Command names that appear within the text of this manual are represented in the following style: "...using the **flush-cache** command clears the bridge cache."

Subsystem names that appear within the text of this manual are represented in the following style: "...to access the **bridge** subsystem..."

Parameter names that appear within the text of this manual are represented in the following style: "...using *<seg-list>* allows you to specify the segments for which you want to display the specified bridge statistics."

Any messages that appear on the screen during software installation and network interface administration are shown in Courier font to distinguish them from the rest of the text as follows:

```
.... Are all four conditions true?
```

Important Information Indicators

To call your attention to safety and otherwise important information that must be reviewed to ensure correct and complete installation, as well as to avoid damage to the FORE Systems product or to your system, FORE Systems utilizes the following *WARNING/CAUTION/NOTE* indicators.

WARNING statements contain information that is critical to the safety of the operator and/or the system. Do not proceed beyond a *WARNING* statement until the indicated conditions are fully understood or met. This information could prevent serious injury to the operator, damage to the FORE Systems product, the system, or currently loaded software, and is indicated as follows:

WARNING!

Hazardous voltages are present. To reduce the risk of electrical shock and danger to personal health, follow the instructions carefully.

CAUTION statements contain information that is important for proper installation/operation. Compliance with **CAUTION** statements can prevent possible equipment damage and/or loss of data and are indicated as follows:

CAUTION

You risk damaging your equipment and/or software if you do not follow these instructions.

NOTE statements contain information that has been found important enough to be called to the special attention of the operator and is set off from the text as follows:

NOTE

If you change the value of the LECS control parameters while the LECS process is running, the new values do not take effect until the LECS process is stopped, and then restarted.

Laser Warning

Every FORE Systems network module having a single mode fiber optic interface contains a Class 1 laser.

> **Class 1 Laser Product:**
> **This product conforms to**
> **applicable requirements of**
> **21 CFR 1040 at the date of**
> **manufacture.**

Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation.

WARNING!

Do not stare into the beam or view the beam with optical instruments.

Safety Precautions

For your protection, observe the following safety precautions when setting up equipment:

- Follow all warnings and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source matches the voltage and frequency inscribed on the equipment's electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous voltages may be present. Conductive foreign objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. FORE Systems, Inc., is not responsible for regulatory compliance of a modified FORE product.

Preface

CHAPTER 1   **Introduction**

This manual contains a detailed description of the AMI `configuration` commands. The main `configuration` menu can be found at the root level. There are several commands available under `configuration`. Commands that are submenus are immediately followed by a ">" symbol. Typing `configuration ?` at the prompt at the root level displays the `configuration` commands as follows:

```
myswitch::> configuration ?
  alarms>          atmarp>          atmroute>        board>
  cec>             ces>             fratm>           funi>
  http>            ilmiproxy>       ip>              lane>
  module>          nsap>            port>            profile>
  qos>             qosext>          rs232>           security>
  snmp>            spans>           spvx>            switch>
  system>          timing>          signalling>      upc>
  vcc>             vpc>             vpt>
```

Each of these commands has a submenu of commands. This manual describes the commands and menus from `configuration alarms>` to `configuration nsap>`. Please see Part 2 of the *AMI Configuration Commands Reference Manual* for descriptions of the commands and menus from `configuration port>` to `configuration vpt>`.

NOTE   Not all of the above commands are displayed on every platform. The `cec>` commands are only displayed on the platforms that can support a CEC-Plus. The `ces>` commands are only displayed on the platforms that can support CEM network modules. The `fratm>` and `funi>` commands are only displayed on the platforms that can support *FramePlus* network modules.

Alarms Configuration Commands

This submenu allows you to configure alarms. You can display the list of available subcommands by typing **alarms ?** at the **configuration** level.

```
myswitch::configuration> alarms ?
  show              enable          disable         reset
  module>
```

2.1   Displaying Alarm Conditions

This command lets you display the status of all alarms. The linkFailed, spansFailed, netmodRemovedHighPrio, and netmodRemovedLowPrio alarms are available on all switches. The powerSupplyInputFailed and tempSensorOverTemp alarms are available on all switches, except an LE 155 and an LE 25. The powerSupplyOutputFailed alarm is only available on an ASX-200BX, an ASX-1000, an ASX-4000, a TNX-210, and a TNX-1100. The fanBankFailed alarm is only available on an ASX-4000, ASX-1000 and a TNX-1100. The powerSupplyOverCurrent and powerSupply5VoltFailed alarms are only available on a DC-powered ASX-1000 and TNX-1100. The fabricRemoved alarm is only available on an ASX-4000. Enter the following parameters:

<div style="float:right">

Alarms Configuration Commands

</div>

```
myswitch::configuration alarms> show
AlarmType                   AlarmStatus   MinorAlarm   MajorAlarm
powerSupplyInputFailed      inactive      disabled     enabled
powerSupplyOutputFailed     inactive      disabled     enabled
fanBankFailed               inactive      disabled     enabled
tempSensorOverTemp          inactive      disabled     enabled
linkFailed                  inactive      enabled      disabled
spansFailed                 inactive      enabled      disabled
powerSupplyOverCurrent      inactive      disabled     enabled
powerSupply5VoltFailed      inactive      disabled     enabled
netmodRemovedHighPrio       inactive      disabled     disabled
netmodRemovedLowPrio        inactive      disabled     disabled
fabricRemoved               inactive      disabled     disabled
Major alarm relay status: off
Minor alarm relay status: off
```

The fields in this display are defined as follows:

Field	Description
AlarmType	The name of the alarm.
AlarmStatus	Shows if the state of the alarm is **active** (alarming) or **inactive** (not alarming). An alarm is active if the underlying condition is detected. For power supplies, the input failed alarm condition is active if the input voltage is not within the nominal range for the supply. This does not necessarily mean that an output failure will result. A power supply output failure condition is active if any power supply is failing or if it is physically removed.
MinorAlarm	**disabled** means that this alarm type will not cause a minor alarm. **enabled** means that this alarm type will cause a minor alarm.
MajorAlarm	**disabled** means that this alarm type will not cause a major alarm. **enabled** means that this alarm type will cause a major alarm.
Major alarm relay status	**off** means no major alarms are currently active. **on** means one or more major alarms are currently active. Look at the AlarmStatus field to see which condition is in a state of alarm.
Minor alarm relay status	**off** means no minor alarms are currently active. **on** means one or more minor alarms are currently active. Look at the AlarmStatus field to see which condition is in a state of alarm.

2.2　Enabling an Alarm

This command lets you enable an alarm. Enter the following parameters:

```
myswitch::configuration alarms> enable (major | minor) <alarm type>
```

These parameters are defined as follows:

Parameter	Description
major	Indicates that the designated alarm type causes a major alarm when that condition occurs.
minor	Indicates that the designated alarm type causes a minor alarm when that condition occurs.
alarm type	Indicates the kind of alarm condition. Valid alarm types are displayed in the AlarmType field for **conf alarms show**.

For example, to enable an overtemperature condition that is detected by the overtemperature sensor as a major alarm, enter the following parameters:

```
myswitch::configuration alarms> enable major tempSensorOverTemp
```

To verify that the change has taken effect, you can display the alarms:

```
myswitch::configuration alarms> show
AlarmType                 AlarmStatus    MinorAlarm    MajorAlarm
powerSupplyInputFailed     inactive       disabled      enabled
powerSupplyOutputFailed    inactive       disabled      enabled
fanBankFailed              inactive       disabled      enabled
tempSensorOverTemp         inactive       disabled      enabled
linkFailed                 inactive       enabled       disabled
spansFailed                inactive       enabled       disabled
netmodRemovedHighPrio      inactive       disabled      disabled
netmodRemovedLowPrio       inactive       disabled      disabled
Major alarm relay status: off
Minor alarm relay status: off
```

Alarms Configuration Commands

2.3   Disabling an Alarm

This command lets you disable an alarm. Enter the following parameters:

```
myswitch::configuration alarms> disable (major | minor) <alarm type>
```

These parameters are defined as follows:

Parameter	Description
major	Indicates that a major alarm is being disabled.
minor	Indicates that a minor alarm is being disabled.
alarm type	Indicates the kind of alarm condition. Valid alarm types are displayed in the `AlarmType` field for **conf alarms show**.

For example, to disable a link failure as a minor alarm, enter the following parameters:

```
myswitch::configuration alarms> disable minor linkFailed
```

To verify that the change has taken effect, you can display the alarms:

```
myswitch::configuration alarms> show

AlarmType                    AlarmStatus    MinorAlarm    MajorAlarm
powerSupplyInputFailed       inactive       disabled      enabled
powerSupplyOutputFailed      inactive       disabled      enabled
fanBankFailed                inactive       disabled      enabled
tempSensorOverTemp           inactive       disabled      enabled
linkFailed                   inactive       disabled      enabled
spansFailed                  inactive       enabled       disabled
netmodRemovedHighPrio        inactive       disabled      disabled
netmodRemovedLowPrio         inactive       disabled      disabled
Major alarm relay status: off
Minor alarm relay status: off
```

2.4   Resetting an Alarm

This command lets you reset either the linkFailed alarm, the spansFailed alarm, or both alarms. Enter the following parameters:

```
myswitch::configuration alarms> reset (<alarm type> | all)
```

These parameters are defined as follows:

Parameter	Description
alarm type	Indicates which alarm to reset. Can be either **linkFailed** or **spansFailed**.
all	Indicates that both the **linkFailed** and the **spansFailed** alarms will be reset.

For example, to reset the linkFailed alarm, enter the following parameters:

```
myswitch::configuration alarms> reset linkFailed
Alarm linkFailed reset.
```

To verify that the change has taken effect, you can display the alarms:

```
myswitch::configuration alarms> show
AlarmType                    AlarmStatus   MinorAlarm   MajorAlarm
powerSupplyInputFailed       inactive      disabled     enabled
powerSupplyOutputFailed      inactive      disabled     enabled
fanBankFailed                inactive      disabled     enabled
tempSensorOverTemp           inactive      disabled     enabled
linkFailed                   inactive      enabled      disabled
spansFailed                  active        enabled      disabled
netmodRemovedHighPrio        inactive      disabled     disabled
netmodRemovedLowPrio         inactive      disabled     disabled
Major alarm relay status: off
Minor alarm relay status: on
```

Alarms Configuration Commands

2.5   Network Module Alarm Commands

These commands, used in conjunction with the **conf alarm enable** and **disable** com-
mands, allow you to configure alarms that notify you of the physical removal of a network
module or a port card from a particular slot in the switch fabric. You can display the list of
available subcommands by typing **module ?** at the alarms level.

```
myswitch::configuration alarms> module ?
  show              assign
```

2.5.1   Displaying the Network Module Alarm Configuration

This command lets you display the current configuration of the network module alarms. Enter
the following parameters:

```
myswitch::configuration alarms module> show
 NetmodSlot    NetmodPrio    NetmodStatus
 A             none          present
 B             high          present
 C             low           present
 D             none          absent
```

The following is displayed on an ASX-4000:

```
 NetmodSlot    NetmodPrio    NetmodStatus
 1A            none          absent
 1B            none          absent
 1C            none          absent
 1D            none          absent
 2A            none          present
 2B            none          present
 2C            none          absent
 2D            none          absent
 3A            none          absent
 3B            none          absent
 3C            none          absent
 3D            none          absent
 4A            none          absent
 4B            none          absent
 4C            none          absent
 4D            none          absent
```

The fields in this display are defined as follows:

Field	Description
NetmodSlot	The network module slot in the switch fabric.
NetmodPrio	The priority assigned to this slot. Can be **none**, **high**, or **low**. **high** means that the `netmodRemovedHighPrio` alarm is activated when a network module or port card is removed from that slot. **low** means that the `netmodRemovedLowPrio` alarm is activated when a network module is removed from that slot. **none** means that no alarm is activated when a network module is removed from that slot. The default is **none**.
NetmodStatus	Shows `present` or `absent` depending on whether or not a network module is currently installed in that slot.

2.5.2   Configuring Network Module Alarm Priorities

This command lets you configure priorities for the network module alarms for each slot in the switch fabric. There are four slots labeled A, B, C, and D. These alarms alert you when a network module has been removed from that slot. On an ASX-4000, the slots are labeled 1A/B, 1C/D, etc. for each port card. Each port card contains two logical network modules. On an ASX-4000, these alarms alert you when a port card has been removed from those slots. Enter the following:

```
myswitch::configuration alarms module> assign <slot> <priority>
```

These parameters are defined as follows:

Parameter	Description
slot	The slot in the switch fabric for which you want to configure an alarm priority. Can be **A**, **B**, **C**, or **D**.
priority	The priority you want to assign to the slot. Each slot in the fabric may be assigned a different priority: high, **low**, or **none**. **high** means that the `netmodRemovedHighPrio` alarm is activated when a network module or port card is removed from that slot. **low** means that the `netmodRemovedLowPrio` alarm is activated when a network module or port card is removed from that slot. **none** means that no alarm is activated when a network module or port card is removed from that slot. The default is **none**.

For example, to configure slot D with a high priority, enter:

```
myswitch::configuration alarms module> assign D high
```

NOTE

On an ASX-4000, you need to include the fabric number in the slot designation. For example, you would enter 2D instead of D to refer to the logical network module on the port card in slot 2C/D.

This assigns a high priority to slot D. Then you need to decide if you want the high priority to be a major or minor alarm (the major and minor alarms for network module removal are disabled by default), and enable it as such, as follows:

```
myswitch::configuration alarms> enable major netmodRemovedHighPrio
```

This makes the high priority a major alarm. If a network module is then removed from slot D, the NetmodStatus field shows absent as follows:

```
myswitch::configuration alarms module> show
 NetmodSlot    NetmodPrio    NetmodStatus
 A             none          present
 B             none          present
 C             low           present
 D             high          absent
```

Also, the AlarmStatus for netmodRemovedHighPrio is active and the Major alarm relay status is on as follows:

```
myswitch::configuration alarms> show
AlarmType                   AlarmStatus    MinorAlarm    MajorAlarm
powerSupplyInputFailed      inactive       disabled      enabled
powerSupplyOutputFailed     inactive       disabled      enabled
fanBankFailed               inactive       disabled      enabled
tempSensorOverTemp          inactive       disabled      enabled
linkFailed                  inactive       enabled       disabled
spansFailed                 inactive       enabled       disabled
netmodRemovedHighPrio       active         disabled      enabled
netmodRemovedLowPrio        inactive       enabled       disabled
Major alarm relay status: on
Minor alarm relay status: off
```

CHAPTER 3 ATMARP Configuration Commands

These commands allow you to manage the ATM ARP (address resolution protocol) features. You can display the list of available subcommands by typing **?** at the **atmarp** level.

```
myswitch::configuration atmarp> ?
arpserver>      delete        flush         getnsap
mapnsap         newclassicalip  newforeip     show
```

3.1   ARP Server Configuration Commands

These commands let you configure the RFC-1577 ATM ARP server. You can get to this level by entering arpserver at the configuration level. By entering **?** at the **arpserver** level, the list of available subcommands for arpserver is displayed.

```
myswitch::configuration atmarp arpserver> ?
  show              set
```

3.1.1   Displaying the ARP Server Address

This command shows which interfaces are enabled to be the RFC-1577 ATM ARP server for the IP network. Enter the following parameters:

```
myswitch::configuration atmarp arpserver> show <interface>
myswitch::configuration atmarp arpserver> show
Interface   ARP Server Addr                                 Enabled
qaa0        0x47.0005.80.ffe100.0000.f21a.3445.0020481a3445.00 Yes
qaa1        0x47.0005.80.ffe100.0000.f21a.3445.0020481a3445.01 Yes
qaa2        0x47.0005.80.ffe100.0000.f21c.06db.0020481c06db.02 No
qaa3        0x47.0005.80.ffe100.0000.f21a.3445.0020481a3445.03 Yes
```

The fields in this display are defined as follows:

Field	Description
Interface	The classical IP interfaces for this switch.
ARP Server Addr	The ARP server address for this interface.
Enabled	Shows whether or not ARP server service is enabled for this interface.

3.1.2   Setting the ARP Server Address

This command allows you to set the address of the RFC-1577 ATM ARP server. Enter the following parameters:

```
myswitch::configuration atmarp arpserver> set <NSAPaddress> [<interface>]
```

These parameters are defined as follows:

Parameter	Description
NSAPAddress	The ATM network layer address for the RFC-1577 ATM ARP server.
interface	Enter the Classical IP interface that will be the arpserver: qaa0, qaa1, qaa2, or qaa3. The default is qaa0.

The switch itself can be used as an ARP server. To do this, set the ARP server address to be the NSAP address of the switch's control port (port CTL).

3.2   Deleting an ARP Entry

This command allows you to remove an ARP entry from the ATM ARP cache. Enter the following parameters:

```
myswitch::configuration atmarp> delete <host>
```

These parameters are defined as follows:

Parameter	Description
host	The IP address of the endstation for which the outgoing ARP entry is to be deleted.

NOTE → If you have ILMI enabled on your switch, ILMI creates an ATM ARP cache entry for each address that it registers. These entries <u>cannot</u> be deleted using this command.

3.3   Flushing the ATM ARP Cache

This command enables you to delete the contents of the ATM ARP cache. Only dynamic ARP cache entries are removed. The switch asks you to verify that flushing the ARP cache is the desired action. Enter the following parameters:

```
myswitch::configuration atmarp> flush
Flush the ATM ARP cache [n]? n
```

To cancel the command, type **n** and press **<ENTER>**, or simply type **<ENTER>**. To flush the ARP cache, type **y** and press **<ENTER>**.

3.4   Getting the NSAP Address for a Classical IP Interface

This command displays the NSAP address of a Classical IP interface. Enter the following:

```
myswitch::configuration atmarp> getnsap [<interface>]
qaa0 NSAP address: 47000580ffe1000000f21a34450020481a344500
```

These parameters are defined as follows:

Parameter	Description
interface	The name of the Classical IP interface to be displayed[1]: qaa0, qaa1, qaa2, or qaa3.

[1.] If no interface is specified, the NSAP address of `qaa0` is displayed.

3.5   Creating an IP to NSAP Address Mapping

This command allows you to create an ATM ARP cache entry mapping a particular IP address to its corresponding NSAP address. Enter the following parameters:

```
myswitch::configuration atmarp> mapnsap <host> <NSAPaddress> [<interface>]
```

These parameters are defined as follows:

Parameter	Description
host	The IP address to be mapped.
NSAPaddress	The NSAP address to be mapped.
interface	The Classical IP interface that should be used to open connections to this NSAP address: qaa0, qaa1, qaa2, or qaa3. The default is qaa0.

3.6   Creating a Classical IP PVC

This command allows you to create a new Classical IP PVC ARP entry. All data is sent LLC/SNAP encapsulated. Enter the following parameters:

```
myswitch::configuration atmarp> newclassicalip <host> <vpi> <vci> [<interface>]
```

These parameters are defined as follows:

Parameter	Description
host	The host IP address of the remote IP endstation.
vpi	The virtual path number of the Classical IP PVC.
vci	The virtual channel number of the Classical IP PVC.
interface	The Classical IP interface to be used for this connection: qaa0, qaa1, qaa2, or qaa3. The default is qaa0.

3.7   Creating a FORE IP PVC ARP Entry

This command enables you to create a FORE IP PVC ARP entry. Data on this PVC is encapsulated using null encapsulation (also known as VC-based multiplexing) as specified in RFC-1483. Enter the following parameters:

```
myswitch::configuration atmarp> newforeip <host> <vpi> <vci> (4|5) [<interface>]
```

These parameters are defined as follows:

Parameter	Description
host	The IP address of the remote host.
vpi	The virtual path number of the FORE IP PVC. Must be 0.
vci	The virtual channel number of the FORE IP PVC.
4 | 5	The connection's ATM Adaptation Layer (AAL) type. The default is 4.
interface	The FORE IP interface to be used for this connection. The default is asx0.

ATMARP Configuration Commands

3.8   Displaying the ATM ARP Entries

This command displays the current ATM ARP cache. Enter the following parameters:

```
myswitch::configuration atmarp> show
myswitch::configuration atmarp>
```

When the prompt is returned with no information displayed, as shown above, then the ATM ARP cache is empty.

The following is an example of an ATM ARP cache.

```
myswitch::configuration atmarp> show
IPaddress       If   VPI  VCI  AAL    Type            Direction
198.29.22.9     asx0 0    63   aal5   foreIpSVC       outgoing
198.29.22.15    asx0 0    231  aal5   foreIpSVC       pending
198.29.22.37    asx0 0    65   aal34  foreIpSVC       pending
IPaddress       If   NSAP Address
198.29.17.3     qaa0 0x47.0005.80.ffe100.0000.f21b.0138.002048102754.00
198.29.17.10    qaa0 0x47.0005.80.ffe100.0000.f21b.0137.002048100be6.00
198.29.17.15    qaa0 0x47.0005.80.ffe100.0000.f21b.0137.00204810048d.00
198.29.17.52    qaa0 0x47.0005.80.ffe100.0000.f21b.0138.0020481b0138.00
```

The fields in this display are defined as follows:

Field	Description
IPaddress	The IP address for this connection.
If	The name of the IP interface for this connection.
VPI	The virtual path number.
VCI	The virtual channel number.
AAL	The AAL type of the given connection.
Type	Shows what kind of connection this is. Can be `foreIpPVC`, `foreIpSVC`, `classicalIpPVC`, or `classicalIpSVC`.
Direction	`Outgoing` means this is an outgoing connection. `Incoming` means this is an incoming connection. `Pending` means that a connection has not (yet) been established. `Incomplete` means that the IP-to-ATM address mapping is not yet known for the given IP address.
NSAP Address	The NSAP address for this connection.

CHAPTER 4 **ATMR Configuration Commands**

These commands let you configure ATM routing on a switch. You can display the list of available subcommands by typing **?** at the **atmroute** level.

```
myswitch::configuration atmroute> ?
  show              domain>         ftpnni>         pnni>
  policy>           spans>
```

4.1   Displaying ATM Routing Information

This command lets you display all of the ATM routing information from various menus in one screen. Enter the following parameters:

```
myswitch::configuration atmroute> show
 Port VPI  Node    Domain SigProto       SigSt  NodeSt HelloSt       PeerSt
 1B1  0    ftpnni  1      FT-PNNI(a)     up     up     N/A           N/A
 1B2  0    1       1      privateUNI(a)  up     up     attempt       N/A
 1B3  0    1       1      privateUNI(a)  up     up     attempt       N/A
 1B4  0    1       1      privateUNI(a)  up     up     attempt       N/A
 1CTL 0    N/A     1      privateUNI(a)  up     N/A    N/A           N/A
```

The fields in this display are defined as follows:

Field	Description
Port	The port through which the link is attached to the node.
VPI	The virtual path number on the port.
Node	For FT-PNNI, shows ftpnni. For PNNI, shows the operational index number of the node for which the ATM routing information is being displayed.
Domain	The index number of the domain to which this node belongs.
SigProto	The signalling protocol being used on this port. **(a)** means that the switch is trying to automatically configure the protocol and version of that protocol that its peers are using.

Field	Description
SigSt	The current state of this interface. If the state is **up**, this interface is operational. This is the normal state for a interface that is connected to another FORE Systems' ATM switch or host. If the state is **down**, this interface is not operational. This can be due to a lack of a physical connection or due to a software problem.
NodeSt	The administrative state of this node. **up** means the node is active. **down** means the node is inactive.
HelloSt	The state of the hello protocol running between the peer nodes. down means that the link is not usable, so no routing packets are sent or received over it. attempt means that either no hellos or hellos with mismatch information have been received from the neighbor, and attempts are being made to reach the neighbor by sending hellos at the specified hello interval. oneWayInside means that hellos have been received from the neighbor and the neighbor has established that they are peers, but the neighbor's remote node ID and remote port ID are 0. twoWayInside means that hellos have been received from the neighbor, the neighbor has established that they are peers, and the neighbor has sent the correct remote node ID and remote port ID. Bi-directional communication can occur over this link. oneWayOutside means that hellos have been received from the neighbor and the neighbor has established that they are from different peer groups, but the neighbor's remote node ID and remote port ID are 0. twoWayOutside means that hellos have been received from the neighbor, the neighbor has established that they are from different peer groups, and the neighbor has sent the correct remote node ID and remote port ID, but the nodal hierarchy list does not include a common peer group. commonOutside means that a common level of the routing hierarchy has been found and bi-directional communication can occur over this link.
PeerSt	N/A means the protocol is something other than PNNI, so this field is not applicable. When the protocol is PNNI, shows the state of the database exchange protocol running between this node and the neighboring peer listed. npdown means there are no active links to the neighboring peer. negotiating means the two peers are deciding which one will start the initial topology database exchange. exchanging means this node is sending its topology database to the neighboring node. loading means this node is receiving the neighboring node's topology database. full means this node has received all PTSEs known to be available from the neighboring peer. Links to the neighboring peer can now be advertised in PTSEs.

You can also display advanced ATM routing information as follows:

```
myswitch::configuration atmroute> show [<port> [<vpi>]] [advanced]
myswitch::configuration atmroute> show advanced
 Port VPI  Node   Domain SigProto      SigSt  NodeSt HelloSt     PeerSt
 1B1  0    ftpnni 1      FT-PNNI(a)    up     up     N/A         N/A
           ConfNode ILMI  ForeLevel    ForeArea    PnniLinkType
           0        up    4            4           N/A
           DomainProto    DomainPrefix
           gateway        0x47.0005.80.ffe100.afce.efed.0000
           FtpnniPgMask    FtpnniPrefix                        Border
           0               0x47.0005.80.ffe100.afce.efed.0000  disabled
 Port VPI  Node   Domain SigProto      SigSt  NodeSt HelloSt     PeerSt
 1B2  0    1      1      privateUNI(a) up     up     attempt     N/A
           ConfNode ILMI  ForeLevel    ForeArea    PnniLinkType
           1        up    5            5           unknown
           DomainProto    DomainPrefix
           gateway        0x47.0005.80.ffe100.afce.efed.0000
           PnniPgId                            Ptses
           80:47.000580ffe100afcefe000000      3
           PnniNodeId
           80:160:47.000580ffe100afceefed0000.ff1a2d0f0001.00
```

The fields in this display are defined as follows:

Field	Description
ConfNode	The PNNI node index number for a user-configured interface that was created through the **conf atmr pnni interface** menu.
ILMI	**Up** means that ILMI is operational for this interface. **Down** means that ILMI is not operational for this interface.
ForeLevel	The level of the area used in FORE's hierarchy support.
ForeArea	The ID of the area in FORE's hierarchy support to which this node belongs.
PnniLinkType	Shows what kind of PTSE this is. Can be nodalInformation, internalReachableAddresses, externalReachableAddresses, nodalStateParameters, horizontalLinks, or uplinks.
DomainProto	The default protocol used on all Network-to-Network (NNI) interfaces in the domain. ftpnni means *ForeThought* PNNI is used for this domain. pnni means ATM Forum PNNI is used for this domain. gateway means this domain includes both the PNNI and FT-PNNI networks so that reachability information can be leaked dynamically between peer groups of both types. This will be a gateway switch (a switch that has one FT-PNNI node and one PNNI node).

ATMR Configuration Commands

Field	Description
DomainPrefix	The 13-byte prefix of the domain used in ILMI registration, in the FT-PNNI switch prefix, and the PNNI node ID. This prefix is also used as the default 13-byte policy summary.
FtpnniPgMask	The mask that gives the number of leading bits in the switch prefix used to aggregate the addresses that belong to this *ForeThought* PNNI peer group. This field is displayed only if the protocol is FT-PNNI.
FtpnniPrefix	The default NSAP prefix for this ATM switch that is used in the ILMI address registration message and in the hello indication FT-PNNI message. This field is displayed only if the protocol is FT-PNNI.
Border	A border switch has a link to at least one other switch that belongs to a different peer group. A border switch advertises reachability information about its peer group to switches outside of its peer group. `enabled` means this switch acts as a *ForeThought* PNNI border switch. `disabled` means this switch does not act as a *ForeThought* PNNI border switch. This field is displayed only if the protocol is FT-PNNI.
PnniPgId	The peer group ID of the peer group to which this node belongs. This field is displayed only if the protocol is PNNI.
Ptses	The number of PTSEs that belong to this node. This field is displayed only if the protocol is PNNI.
PnniNodeId	The ID of this node. This field is displayed only if the protocol is PNNI.

The first nine fields in this display are defined in the same manner as those listed in the previous **show** command.

4.2   Configuring a Domain

These commands allow you to configure a domain on a switch. An ATM routing domain is a group of areas that are configured to dynamically exchange reachability information with one another. This allows connectivity between end systems belonging to different areas. Reachability information is exchanged between domains through static configuration only. A switch may be part of multiple domains and there may be multiple domains configured on a switch. (For more information about domains and areas, see the *ATM Switch Network Configuration Manual*.) You can display the available subcommands by typing **?** at the `domain` level.

```
myswitch::configuration atmroute domain> ?
  delete          modify          new             show
```

4.2.1   Deleting a Domain

This command allows you to delete an existing domain. Enter the following parameters:

```
myswitch::configuration atmroute domain> delete <domain ID>
```

4.2.2   Modifying a Domain

This command allows you to modify an existing domain. Enter the following parameters:

```
myswitch::configuration atmroute domain> modify <domain ID> [-name <name>]
[<ftpnni | pnni | gateway>][-prefix <prefix>] [-defsum <enable | disable>]
```

NOTE

If you modify the domain, the switch prompts for a reboot. You must type **y** and reboot the switch for your change to take effect.

NOTE

If the existing domain is pnni or gateway, and the prefix of the domain is changed, you must first disable the PNNI node using **conf atmr pnni node admin** *<nodeid>* **down**. This removes any transient routes related to the node address with the old prefix.

ATMR Configuration Commands

4.2.3   Creating a Domain

This command allows you to create a new domain. Enter the following parameters:

```
myswitch::configuration atmroute domain> new <domain ID> [-name <name>]
[-prefix <prefix>] [-defsum <enable | disable>]
```

The parameters for delete, modify, and new are defined as follows:

Parameter	Description
domain ID	The index number of this domain. Displayed in the the `Index` field under **conf atmroute domain show**.
name	The optional, user-assigned name of this domain. Displayed in the `Name` field under **conf atmroute domain show**.
ftpnni pnni gateway	The default routing protocol to use on all Network-to-Network (NNI) interfaces in the domain. `ftpnni` means *ForeThought* PNNI is used for this domain. `pnni` means ATM Forum PNNI is used for this domain. `gateway` means this domain includes both the PNNI and FT-PNNI networks so that reachability information can be leaked dynamically between peer groups of both types. This will be a gateway switch (a switch that has one FT-PNNI node and one PNNI node). The default routing protocol is `ftpnni`.
prefix	The 13-byte prefix of the domain used in ILMI registration, used in the FT-PNNI switch prefix, and used in determining the PNNI node ID, PNNI peer group, and PNNI ATM address. This prefix is also used as the default 13-byte policy summary. If you have two PNNI domains in a switch, be sure to assign each one a different prefix.
-defsum <enable | disable>	Enables or disables the ability to summarize all local addresses into a default, 10-byte summary switch prefix, and advertise this summary in PNNI as an internal reachable address PTSE and in FT-PNNI topology. Modifying this parameter does not require a reboot; the change takes effect immediately. The default value is `disable`.

4.2.4   Displaying Domain Information

This command allows you to display information about all of the existing domains. Enter the
following parameters:

```
myswitch::configuration atmroute domain> show
ID  Name              Defproto  Prefix                        Defsum
1   default           ftpnni    0x47.0005.80.ffe100.0000.f21a.3445   ENABLED
```

The fields in this display are defined as follows:

Field	Description
ID	The index number of this domain.
Name	The user-assigned name of this domain.
Defproto	The default routing protocol used on all Network-to-Network (NNI) interfaces in the domain. ftpnni means *ForeThought* PNNI is used for this switch in this domain. pnni means ATM Forum PNNI is used for this switch in this domain. gateway means this switch includes both the PNNI and FT-PNNI networks in this domain so that reachability information can be leaked dynamically between peer groups of both types.
Prefix	The address prefix of the domain.
Defsum	Shows if the ability to summarize all local addresses into a default, 10-byte summary switch prefix is enabled or disabled.

4.3 *ForeThought* PNNI Configuration Commands

These commands allow you to configure *ForeThought* PNNI on a switch. You can display the list of available subcommands by typing **?** at the **ftpnni** level.

```
myswitch::configuration atmroute ftpnni> ?
  lbubr            border            dtl>              forearea
  forelevel        hello             maxhop            metric>
  minthresh        nsapindication    pgmask            pgsncost
  prefix           propmult          staticroute>      swmask
  vcmark           show
```

4.3.1 Configuring Load-Balanced UBR Routing

Load-balanced UBR (LBUBR) routing lets you distribute UBR SPVCs evenly across links with varying Available Cell Rate capacities. For example, currently, if a UBR call from one switch to another can go through three different equal cost paths, it always chooses the highest bandwidth, by default. This saturates one path, and leaves the two alternate paths unused. Instead, by enabling this feature, UBR calls are distributed evenly based on the estimated bandwidth of individual calls and on the percentage of usage among all available links, such that the percentage of available bandwidth on any given link is maximized. (However, when paths are overbooked, normal load balancing is done. In this case, the alternate paths are equally preferable and routing is not done according to weighted available bandwidth.)

You can assign the estimated bandwidth value, in cells per second, using the **-estbw** *<kcps>* option under **conf upc new**, and then apply it to UBR SPVCs using the **-fupc** *<index>* and **-bupc** *<index>* options under **conf spvx spvcc pnni new**. You can also assign a default estimated UBR bandwidth value, in cells per second, using the **-defaultbw** *<kcps>* option under **conf spvx spvcc pnni parameters lbubr**. See Part 2 of this manual for more information about each of these options.

This estimated bandwidth value is then used by routing to determine if any particular link has the available bandwidth to support the connection, and to provide an estimate of traffic flow through any link in the network. This information is used by path computation when calculating routes. This command allows you to enable or disable load-balanced UBR routing on a switch. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> lbubr (enable | disable)
```

These parameters are defined as follows:

Parameter	Description
enable | disable	**enable** means load-balanced UBR routing is used on this switch. UBR SPVCs are distributed evenly based on the estimated bandwidth of individual calls and on the percentage of usage among all available links. **disable** means load-balanced UBR routing is not used on this switch. The default is **disable**.

4.3.2   Changing the FT-PNNI Border Switch Functionality

A switch that has a link to another switch that belongs to a different peer group is considered a border switch. A border switch advertises reachability information about its peer group to switches outside of its peer group. You should enable border switch functionality on all switches in a peer group that have direct outside links to other peer groups. This command lets you designate whether or not this switch will act as a *ForeThought* PNNI border switch. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> border (enable | disable)
```

> **NOTE** The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

These parameters are defined as follows:

Parameter	Description
enable	Entering enable (and rebooting) means that this switch will act as a *ForeThought* PNNI border switch.
disable	Entering disable (and rebooting) means that this switch will not act as a *ForeThought* PNNI border switch.

4.3.3 DTL Configuration Commands

These commands let you create, delete, and display Designated Transit Lists (DTLs). A DTL is a source route for the *ForeThought* PNNI (FT-PNNI) router which specifies the preferred call routing for the SVC portion of a directed SPVC. A DTL is a source route (index) and each entry (row) in the DTL represents a single hop in the source route. Each hop is represented by a FT-PNNI node, and the output port and VPI at that node. The FT-PNNI node's address is determined by the switch's NSAP prefix and mask. You can display the list of available sub-commands by typing **?** at the **dtl** level.

```
myswitch::configuration atmroute ftpnni dtl> ?
  delete          dtlhopid        dtlhopidparse      modify
  new             show
```

4.3.3.1 Deleting a DTL Entry

This command enables you to remove an existing DTL. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> delete <index> (<row>|all)
```

For example, if you want to delete row 2 of index entry 9, enter the following:

```
myswitch::configuration atmroute ftpnni dtl> delete 9 2
```

Similarly, if you want to delete every row in index entry 9, enter the following:

```
myswitch::configuration atmroute ftpnni dtl> delete 9 all
```

These parameters are defined as follows:

Parameter	Description
index	The index number of a DTL. Displayed in the Index field under **conf atmroute ftpnni dtl show**.
row	The row number of the individual entry within a given DTL. An entry in the DTL is given a row number equal to its position in the source route. Thus, the entry corresponding to the first hop is row 1, the second hop is row 2, and so on. Displayed in the Row field under **conf atmroute ftpnni dtl show**.
all	Indicates that you want to delete all of the entries within a given DTL.

4.3.3.2 Computing a DTL Hop ID

This command lets you compute a DTL hop ID for a switch. This command displays the pre-fix, mask, port and vpi information, which is encoded in a hex string. This information can be used as the *<dtlhopid>* parameter under **conf atmr ftpnni new**. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> dtlhopid <port> <vpi>
```

These parameters are defined as follows:

Parameter	Description
port	The output port number, in BNP notation, for this DTL entry. In *ForeThought* PNNI, the node IDs are prefixes that represent both the ID of the node and a summary of reachable addresses. This is also the port to which the next node in the DTL (if there is one) is connected.
vpi	The output path number in the above output port. This is the signalling path to the peer given by the next node in the DTL.

For example:

```
myswitch::configuration atmroute ftpnni dtl> dtlhopid 1c1 0
DTL Id : 0x47.0005.80.ffe100.0000.f21a.3bae.0068.0010.0000
```

4.3.3.3 Parsing a DTL Hop ID

When you enter a DTL hop ID for a switch, this command displays the prefix, mask, port, and vpi information that are encoded in the ID. This command also checks to see if the ID belongs to this switch. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> dtlhopidparse <dtlhopid>
```

This parameter is defined as follows:

Parameter	Description
dtlhopid	The DTL hop ID for a switch, in hex string format, which contains the encoded prefix, mask, port, and vpi information. Use the **conf atmr ftpnni dtl dtlhopid** command to obtain this ID.

For example:

```
myswitch::configuration atmroute ftpnni dtl> dtlhopidparse
0x47.0005.80.ffe100.0000.f21a.3bae.0068.0010.0000
NSAP prefix                        Mask  Port  VPI
0x47.0005.80.ffe100.0000.f21a.3bae   104  1C1   0
```

The fields in this display are defined as follows:

Field	Description
NSAP prefix	The NSAP prefix part of the node ID of this DTL entry. In *ForeThought* PNNI, the node IDs are prefixes that represent both the ID of the node and a summary of reachable addresses.
Mask	The mask corresponding to the prefix, which gives the length of the above prefix in number of bits.
Port	The output port for the node given by the above node ID. This is also the port to which the next node in the DTL (if there is one) is connected.
VPI	The output path in the above output port. This is the signalling path to the peer given by the next node in the DTL.

If the ID does not belong to this switch, you are prompted if you want to continue anyway.

```
myswitch::configuration atmroute ftpnni dtl> dtlhopidparse
0x47.0005.80.ffe100.0000.f31b.3bff.0025.0067.0000

The dtlhopid does not belong to this switch
Do you want to continue parsing [n]? y
NSAP prefix                        Mask  Port  VPI
0x47.0005.80.ffe100.0000.f31b.3bff   104  1D3   0
```

4.3.3.4   Modifying a DTL Entry

This command allows you to modify a DTL entry. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> modify <index> <row>\
(prefix | mask | port | vpi |dtlhopid) <new_value>
```

For example, to modify the port of index entry 9, row 2 to port D1, enter the following:

```
myswitch::configuration atmroute ftpnni dtl> modify 9 2 port D1
```

Similarly, to modify the mask of index entry 12, row 1 to a mask of 104, enter the following:

```
myswitch::configuration atmroute ftpnni dtl> modify 12 1 mask 104
```

The parameters for modify and new are defined as follows:
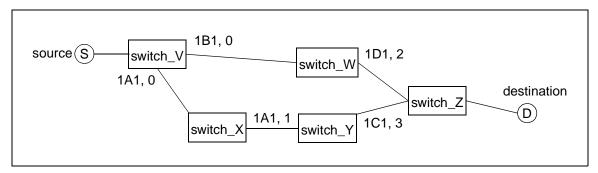
Parameter	Description
index	The index number of a DTL. Displayed in the `Index` field under **conf atmroute ftpnni dtl show**.
row	The row number of the individual entry within a given DTL. An entry in the DTL is given a row number equal to its position in the source route. Thus, the entry corresponding to the first hop is row 1, the second hop is row 2, and so on. Displayed in the `Row` field under **conf atmroute ftpnni dtl show**.
prefix	Indicates the 13-byte NSAP prefix of the node ID of this DTL entry. In *ForeThought* PNNI, the node IDs are prefixes that represent both the ID of the node and a summary of reachable addresses.
mask	The subnet mask is the significant part of the prefix when matching an address. Gives the length of the above prefix in number of bits. Can vary from 0 - 104 (13 bytes).
port	The output port number for the node given by the above node ID. This is also the port to which the next node in the DTL (if there is one) is connected. You only need to specify the network module and port numbers, not the board number (c1 instead of 1c1).
vpi	The output path number in the above output port. This is the signalling path to the peer given by the next node in the DTL.
dtlhopid	The DTL hop ID for a switch that contains encoded prefix, mask, port, and vpi information in a hex string. Use the **conf atmr ftpnni dtl dtlhopid** command to obtain this ID.
new_value	Used only with the **modify** command. Indicates the actual value that you wish to use for one of the above parameters. Can be the new prefix, mask, port, vpi, or dtlhopid.

ATMR Configuration Commands

4.3.3.5 Creating a DTL Entry

This command allows you to create a new DTL. There are two methods for creating DTLs. The first method is the newer and preferred method. This method must be used if there are any ASX-4000 switches in the route. If this method is used, all switches must be running *ForeThought* 5.3.x or later. The second method can only be used for DTLs that do not include any ASX-4000 switches in the route. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> new <index> <row>
(-dtlhopid <dtlhopid> | [<prefix> <mask> <port> <vpi>])
```

One example for each method is given to show how to create DTLs for this simple network.

Method 1

This example, using the newer method, assumes that all of the switches in the DTL path are running *ForeThought* 5.3.x or later.

1. To create a DTL from point S to point D, enter the **conf atmr ftpnni dtlhopid** command on switch_V, switch_X, and switch_Y to obtain the dtlhopid of each switch.

```
switch_V::configuration atmroute ftpnni dtl> dtlhopid 1a1 0
DTL Id : 0x47.0005.80.ffe100.0000.f21a.3bae.0068.0010.0000

switch_X::configuration atmroute ftpnni dtl> dtlhopid 1a1 1
DTL Id : 0x47.0005.80.ffe100.0000.f21a.33af.0075.0010.0000

switch_Y::configuration atmroute ftpnni dtl> dtlhopid 1c1 3
DTL Id : 0x47.0005.80.ffe100.0000.f21a.5bbe.0012.0010.0000
```

2. Use the `conf atmr ftpnni dtl new` command and cut and paste the dtlhopid for each switch.

```
switch_V::configuration atmroute ftpnni dtl> new 1 1
0x47.0005.80.ffe100.0000.f21a.3bae.0068.0010.0000

switch_V::configuration atmroute ftpnni dtl> new 1 2
0x47.0005.80.ffe100.0000.f21a.33af.0075.0010.0000

switch_V::configuration atmroute ftpnni dtl> new 1 3
0x47.0005.80.ffe100.0000.f21a.5bbe.0012.0010.0000
```

This creates a DTL that routes through switch_V, switch_X, switch_Y, and switch_Z. The prefix in each line changes to match the prefix of each switch through which the route is going. Notice that you do not specify switch_Z in the DTL path.

Method 2

This example can only be used if none of the switches in the DTL path is an ASX-4000. When using this older method, you <u>must</u> specify all of the following parameters: `<prefix>`, `<mask>`, `<port>`, and `<vpi>`.

1. To create a DTL from point S to point D, enter something similar to the following:

```
switch_V::configuration atmroute ftpnni dtl> new 1 1
0x47000580ffe1000000f21b19c3 104  1A1 0

This option is meant only for pre FT5.3 transit switches
Are you sure you want to continue [n]? y

switch_V::configuration atmroute ftpnni dtl> new 1 2
0x47000580ffe1000000f21b24a2 104  1A1 1

This option is meant only for pre FT5.3 transit switches
Are you sure you want to continue [n]? y

switch_V::configuration atmroute ftpnni dtl> new 1 3
0x47000580ffe1000000f21b18a1 104  1C1 3

This option is meant only for pre FT5.3 transit switches
Are you sure you want to continue [n]? y
```

This creates a DTL path that routes through switch_V, switch_X, switch_Y, and switch_Z. The prefix in each line changes to match the prefix of each switch through which the route is going. Notice that you do not specify switch_Z in the DTL path.

ATMR Configuration Commands

4.3.3.6 Displaying DTL Hop IDs

This command lets you display the current DTL hop IDs. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> show
Index    Row    DTL Id
   1      1     0x47.0005.80.ffe100.0000.f21a.3bae.0068.0013.0000
          2     0x47.0005.80.ffe100.0000.f21a.3cae.0068.0011.0000
          3     0x47.0005.80.ffe100.0000.f21a.3cae.0068.0001.0000
```

The fields in this display are defined as follows:

Field	Description
Index	The index number of each of the current DTLs.
Row	The row number of each entry within each DTL. Each entry in the DTL has a row number equal to its position in the source route, so the entry corresponding to the first hop is row 1, the second hop is row 2, and so on.
DTL Id	The DTL hop ID, which contains encoded prefix, mask, port, and vpi information in a hex string format.

NOTE
The NSAP prefix, mask, port, and vpi information for the ID can be displayed on the switch to which it belongs using the **conf atmr ftpnni dtlhopidparse** command.

You can also display an individual DTL by entering the following parameters:

```
myswitch::configuration atmroute ftpnni dtl> show [<index>]
myswitch::configuration atmroute ftpnni dtl> show 1
Index    Row    DTL Id
   1      1     0x47.0005.80.ffe100.0000.f21a.3bae.0068.0013.0000
          2     0x47.0005.80.ffe100.0000.f21a.3cae.0068.0011.0000
          3     0x47.0005.80.ffe100.0000.f21a.3cae.0068.0001.0000
```

These fields are defined in the same manner as in the previous example.

If no DTLs exist, the following is displayed:

```
myswitch::configuration nsap dtl> show
No DTLs available
```

4.3.4   Configuring the FORE Area

This command lets you set the value of the FORE Area ID. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> forearea <forearea>
```

The switch cautions you that a reboot is necessary. It asks if you want to continue with the change in case you made a mistake. To abort the change, type **n** or press **<ENTER>**. No change will be made. To continue with the change, type **y**. If you type **y**, you will be asked if you want to reboot the switch. You must reboot the switch for this change to take effect, so type **y** or press **<ENTER>**.

```
This change requires a reboot to prevent any potential routing problem.
Do you want to continue with the change [n]? y
Reboot the switch [y]? y
```

This parameter is defined as follows:

Parameter	Description
forearea	The ID of the area in the FORE hierarchy to which this *ForeThought* PNNI node belongs. This can be a value between 1 and 127. The default is 4.

4.3.5   Configuring the FORE Level

This command lets you set the value of the FORE Level. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> forelevel <forelevel>
```

The switch cautions you that a reboot is necessary. It asks if you want to continue with the change in case you made a mistake. To abort the change, type **n** or press **<ENTER>**. No change will be made. To continue with the change, type **y**. If you type **y**, you will be asked if you want to reboot the switch. You must reboot the switch for this change to take effect, so type **y** or press **<ENTER>**.

```
This change requires a reboot to prevent any potential routing problem.
Do you want to continue with the change [n]? y
Reboot the switch [y]? y
```

This parameter is defined as follows:

Parameter	Description
forelevel	The level of the area to which this *ForeThought* PNNI node belongs in the FORE hierarchy. This can be a value between 1 and 127. The default is 4.

**ATMR Configuration
Commands**

4.3.6   Setting the Hello Indication Interval

Hello indication messages are the "keep alive" messages that two switches send to one another to verify their existence. This command lets you change the interval for *ForeThought* PNNI hello indication messages. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> hello <msec>
```

This parameter is defined as follows:

Parameter	Description
hello	The period of time between transmissions of hello indication messages. The default 500 microseconds.

4.3.7   Setting the Maximum Hop Count

This command lets you set the maximum hop count for the NSAP router. By setting a maximum hop count, you tell the switch to consider only those paths that have less than or equal to the number of hops specified when setting up a connection. If a connection is routed using a path with a large hop count, there is a greater chance that the connection may experience congestion and be delayed or discarded. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> maxhop <hops>
```

This parameter is defined as follows:

Parameter	Description
maxhop	The maximum number of hops to use when routing a connection for the NSAP router. The default is 20 hops.

4.3.8   *ForeThought* **PNNI Metric Configuration Commands**

These commands allow you to configure *ForeThought* PNNI metric sets. You can display the list of available subcommands by typing **?** at the **metric** level.

```
myswitch::configuration atmroute ftpnni> metric ?
  delete            modify            new            show
```

4.3.8.1   Deleting a Metric Set

This command lets you delete a metric set. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni metric> delete <tag>
```

4.3.8.2   Modifying a Metric Set

This command lets you modify a metric set. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni metric> modify <tag> [-cost <cost>]
[-cbrcap <cbrcap>] [-vbrcap <vbrcap>]
```

4.3.8.3   Creating a Metric Set

This command lets you create a metric set that is used when creating a policy for a FT-PNNI static route. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni metric> new <tag> [-cost <cost>]
[-cbrcap <cbrcap>][-vbrcap <vbrcap>]
```

The parameters for delete, modify, and new are defined as follows:

Parameter	Description
tag	Indicates a unique integer that identifies this metric set.
-cost <cost>	The cost of reaching the address encompassed by this policy.
-cbrcap <cbrcap>	The CBR capacity of the link to reach the address encompassed by this policy.
-vbrcap <vbrcap>	The VBR capacity of the link to reach the address encompassed by this policy.

ATMR Configuration Commands

4.3.8.4 Displaying Metric Set Information

This command lets you display metric set information. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni metric> show
  Tag  Cost   CBRCAP   VBRCAP
  1    10     10       10
  Tag  Cost   CBRCAP   VBRCAP
  2    128     10       200
```

The fields in this display are defined as follows:

Field	Description
Tag	The unique integer number that identifies this metric set.
Cost	The cost of reaching the address encompassed by this policy.
CBRCAP	The CBR capacity of the link to reach the address encompassed by this policy.
VBRCAP	The VBR capacity of the link to reach the address encompassed by this policy.

You can also display information about a specific tag as follows:

```
myswitch::configuration atmroute ftpnni metric> show [<tag>]
myswitch::configuration atmroute ftpnni metric> show 2
  Tag  Cost   CBRCAP   VBRCAP
  2    128     10       200
```

If no metrics have been configured, then the following is displayed:

```
myswitch::configuration atmroute ftpnni metric> show
No metric information is available
```

4.3.9 Setting a Minimum Threshold for NSAP Updates

The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which the product of the ACR and the proportional multiplier is a very small number. The minimum threshold is used to prevent excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> minthresh <minthresh>
```

This parameter is defined as follows:

Parameter	Description
minthresh	The minimum threshold bandwidth value for triggering NSAP updates. The default is 50 Kbps.

4.3.10 Setting the NSAP Indication Interval

NSAP indication messages are those messages that update topology information between any two switches. This command lets you set the interval between *ForeThought* PNNI NSAP indication messages. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> nsapindication <msec>
```

This parameter is defined as follows:

Parameter	Description
nsapindication	The period of time between transmissions of NSAP indication messages. The default is 10,000 microseconds.

ATMR Configuration Commands

4.3.11 Setting the *ForeThought* PNNI Peer Group Mask

A peer group mask is the length (in the number of bits) of the peer group ID of a switch. This command enables you to set the *ForeThought* PNNI peer group mask value. This value should be the same for all members of a peer group. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> pgmask <mask>
```

This parameter is defined as follows:

Parameter	Description
pgmask[1]	The mask that gives the number of leading bits in the switch prefix used to aggregate the addresses that belong to this *ForeThought* PNNI peer group. The default is 0.

[1]. The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

4.3.12 Selecting the Method for Computing the Cost of a Link

This command lets you select the method of computing the cost of a link from a border node to the peer group summary node (PGSN). Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> pgsncost (default | user -cost <cost>)
```

These parameters are defined as follows:

Parameter	Description
default	The border node automatically calculates the cost to the PGSN by taking half of the average cost from this node to all other border nodes in this peer group. This cost is dynamic.
user	The link to the PGSN from this border node will be advertised using the value that you specify with the -cost parameter.
-cost <cost>	Enter the administrative weight for the link that you want the border node to use during path computation. The route which takes the least cost is chosen. The default cost for all links in the network is 100.

4.3.13 Setting the *ForeThought* PNNI Switch Prefix

When using *ForeThought* PNNI, a switch fabric is identified by an NSAP switch prefix which consists of 13 fixed bytes. The variable 13-byte mask configured using **conf atmroute ftpnni swmask** determines which bytes are actually significant. This command lets you set the *ForeThought* PNNI prefix on the switch. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> prefix <prefix>
```

This parameter is defined as follows:

Parameter	Description
prefix[1]	The FT-PNNI prefix for this ATM switch that is used in the hello indication FT-PNNI message.

[1] The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

4.3.14 Setting the Proportional Multiplier

This command enables you to set the proportional multiplier for the NSAP router. The proportional multiplier is expressed as a percentage of Available Cell Rate (ACR) on any given link in the network. If the change in percentage of the ACR on any given link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant. The topology tables are updated accordingly for that link. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> propmult <percentage>
```

This parameter is defined as follows:

Parameter	Description
propmult[1]	The threshold above which you consider the change in ACR on any link to be significant. The default is 20%.

[1] If you modify this value, you should modify it on <u>all</u> switches in the network.

4.3.15 Static Route Configuration Commands

These commands let you create, delete, and display static routes. You can display the list of available subcommands by typing **?** at the **staticroute** level.

```
myswitch::configuration atmroute ftpnni staticroute> ?
  delete           new             show
```

4.3.15.1 Deleting a Static Route

This command enables you to remove an existing static route. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni staticroute> delete <NSAP> <mask> <port> <vpi>
```

4.3.15.2 Creating a Static Route

This command allows you to create a static route. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni staticroute> new <NSAP> <mask> <port> <vpi>
[-cost <cost>][-cbr_cap <cbr_cap>] [-vbr_cap <vbr_cap>]
```

The following is an example of how to create an NSAP static route:

```
myswitch::configuration atmroute ftpnni staticroute> new
0x47.0005.80.ffe100.0000.f215.11f2.002048100464.00 152 1c2 0 -cost 200 -cbr_cap 20000
-vbr_cap 30000
```

The parameters for delete and new are defined as follows:

Parameter	Description
NSAP	The complete 20-byte NSAP route address in hexadecimal format.
mask	The bit mask indicating number of high-order bits to use for routing purposes. The default mask for the route to the host is 152 and the default mask for the route to the switch is 104.
port	The port number through which this static route can be reached.
vpi	The UNI signalling path through which this static route can be reached.
-cost <cost>	Used only with the `new` command. The routing metric for this link. There is a cost for each link in a route. The sum of these link costs determines the overall cost of a route. To expedite traffic on a route, try to minimize the overall cost of a route. For a critical route, then, choose a small cost value. For a lesser important route, choose a higher cost value.
-cbr_cap <cbr_cap>	Used only with the `new` command. The maximum CBR (Constant Bit Rate) capacity allowed for any single connection on this route. This number is limited by the actual CBR capacity available on the output link specified for this route.
-vbr_cap <vbr_cap>	Used only with the `new` command. The maximum VBR (Variable Bit Rate) capacity allowed for any single connection on this route. This number is limited by the actual VBR capacity available on the output link specified for this route.

4.3.15.3 Displaying Static Routes

This command lets you display the current static routes. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni staticroute> show
NSAP-address                             Mask Port VPI Cost CBR   VBR
                                                            Mbs   Mbs
47000580ffe1000000f21511f200204810046400 152  1C2  0   200  20.0  30.0
47000580ffe1000000f21511f20020481ee00000 144  1C3  0   100  70.0  60.0
47000580ffe1000000f21511f20020481ff00000 144  1C1  0   100  INF   INF
47000580ffe1000000f21511f20020481ff12300 152  1C3  0   100  INF   INF
```

The fields in this display are defined as follows:

Field	Description
NSAP-address	Shows the 20-byte address for which the static route is configured.
Mask	The bit mask indicating number of high-order bits to use for routing purposes. The default mask for a static route to a host is `152` and the default mask to another switch is `104`.
Port	The port number on which the NSAP route exists.
VPI	The number of the virtual path on which the NSAP static route exists.
Cost	The routing metric for this link. There is a cost for each link in a route. The sum of these costs determines the overall cost of a route. To expedite traffic on a route, try to minimize the overall cost of a route. A small cost value is assigned to a critical route, while a higher cost value is assigned to a lesser important route. The default is `100`.
CBR	The maximum CBR capacity allowed for any single connection on this route. `INF` means that you did not specify a value for this parameter when the route was created, so the value defaults to the capacity available on the outgoing link.
VBR	The maximum VBR capacity allowed for any single connection on this route. `INF` means that you did not specify a value for this parameter when the route was created, so the value defaults to the capacity available on the outgoing link.

You can also display static route information for an individual NSAP address, or NSAP address <u>and</u> mask as follows:

```
myswitch::configuration atmroute ftpnni staticroute> show [<NSAP> [<mask>]]
myswitch::configuration atmroute ftpnni staticroute> show
47000580ffe1000000f21511f200204810046400 152
NSAP-address                             Mask Port VPI Cost CBR   VBR
                                                            Mbs   Mbs
47000580ffe1000000f21511f200204810046400 152  1C2  0   200  20.0  30.0
```

If you have not configured any static routes, then the following message is displayed:

```
myswitch::configuration atmroute ftpnni staticroute> show
No NSAP static route information is available
```

4.3.16 Setting the *ForeThought* PNNI Switch Prefix Mask

This command allows you to select the *ForeThought* PNNI switch prefix mask value. Enter the following parameters:

```
myswitch::configuration atmroute ftpnni> swmask <mask>
```

This parameter is defined as follows:

Parameter	Description
swmask[1]	The mask that gives the number of leading bits in the switch prefix used to aggregate the addresses that belong to the switch in *ForeThought* PNNI. The default is 104.

[1.] The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

4.3.17 Setting a Minimum Virtual Channel Mark

When the number of available virtual channels on a path drops to zero, a link state update is sent out to advertise that there are no more VCs available for use on this path. When the number of VCs indicated by the **vcmark** is available for use on this path again, another link state update is sent out to advertise that there are VCs available for use on this path once again. This command lets you set the **vcmark**, which is the minimum number of virtual channels that need to be to available on a path to make that path usable again. Enter the following:

```
myswitch::configuration atmroute ftpnni> vcmark <vcmark>
```

This parameter is defined as follows:

Parameter	Description
vcmark	The minimum number of virtual channels that need to be available on a path to make that path usable. The default is 20.

4.3.18 Displaying *ForeThought* PNNI Parameters

This command lets you display all of the *ForeThought* PNNI topology parameters. Enter the following:

```
myswitch::configuration atmroute ftpnni> show

Switch NSAP prefix                    0x47.0005.80.ffe100.0000.f21c.078e
Switch Prefix Mask                    104
Peer Group Mask                       0

Hello Indication Interval             500 msec
NSAP Indication Interval              10000 msec
Max hop count for NSAP router         20 hops
Proportional Multiplier               20 %
Minimum Threshold for NSAP updates    50 Kbps
Minimum VC level                      20
Fore Area                             4
Fore Level                            4
Cost of link to PGSN                  100
Cost of link to PGSN computing method default

Load Balanced UBR Routing is disabled
FORE PNNI border switch functionality is enabled
```

The fields in this display are defined as follows:

Field	Description
Switch NSAP prefix	The switch's NSAP prefix.
Switch Prefix Mask	The switch prefix mask value of high-order bits to use for aggregating addresses on the switch for routing purposes.
Peer Group Mask	The peer group mask value of high-order bits to use for aggregating addresses on the switch for routing purposes.
Hello Indication Interval	The period of time between transmissions of hello indication messages.
NSAP Indication Interval	The period of time between transmissions of NSAP indication messages.
Max hop count for NSAP router	The maximum number of hops to use when routing a connection for the NSAP router.
Proportional Multiplier	The threshold above which the change in ACR on any link is considered to be significant.
Minimum Threshold for NSAP updates	The minimum threshold bandwidth value for triggering NSAP updates.

Field	Description
Minimum VC level	The minimum number of VCs that need to be available on a path to make that path usable again after the number of available VCs has dropped to 0.
Fore Area	The ID of the area in the FORE hierarchy to which this *ForeThought* PNNI node belongs.
Fore Level	The level of the area to which this *ForeThought* PNNI node belongs in the FORE hierarchy.
Cost of Link to PGSN	The administrative weight for the link that the border node uses during path computation to the PGSN. The route which takes the least cost is chosen. This field is displayed on border switches only.
Cost of Link to PGSN computing method	The method used for computing the link cost from the border switch to the PGSN. This field is displayed on border switches only. Can be `default` or `user`.
Load Balanced UBR Routing is disabled	If this functionality is `enabled`, load-balanced UBR routing occurs on this switch. UBR SPVCs are distributed evenly based on the estimated bandwidth of individual calls and on the percentage of usage among all available links. If this functionality is `disabled`, load-balanced UBR routing does not occur on this switch.
FORE PNNI border switch functionality is disabled	If this functionality is `enabled`, this switch acts as a *ForeThought* PNNI border switch. If this functionality is `disabled`, this switch does not act as a *ForeThought* PNNI border switch.

ATMR Configuration Commands

4.4   ATM Forum PNNI Configuration Commands

These commands allow you to configure ATM Forum PNNI on a switch. You can display the list of available subcommands by typing **?** at the **pnni** level.

```
myswitch::configuration atmroute pnni> ?
   address>        crankback>      interface>      metric>
   node>           profile>
```

4.4.1   PNNI Address Configuration Commands

These commands allow you to configure ATM Forum PNNI exterior reachability addresses. You can display the list of available subcommands by typing **?** at the **address** level.

```
myswitch::configuration atmroute pnni address> ?
   delete          modify          new            show
```

4.4.1.1   Deleting a PNNI Exterior Reachability Address

This command deletes the PNNI exterior reachability addresses and static routes between two domains called interdomain routes. Enter the following parameters:

```
myswitch::configuration atmroute pnni address> delete pnni <nodeix> <address> <len>
<index> OR/interdomain <domain_id> <address> <len> <dest_domain_id>
```

4.4.1.2   Modifying a PNNI Exterior Reachability Address

This command lets you modify the PNNI exterior reachability address between two domains. Enter the following parameters:

```
myswitch::configuration atmroute pnni address> modify <nodeix> <address> <len> <index>
[-port <port> -vpi <vpi>] [-mtag <mtag>]
```

4.4.1.3 Creating a PNNI Exterior Reachability Address

This command lets you create a PNNI exterior reachability address between two domains so they can share reachability information. Enter the following parameters:

```
myswitch::configuration atmroute pnni address> new pnni <nodeix> <address> <len> <index>
-port <port> -vpi <vpi> [-mtag <mtag>] [-vpcap (true|false)] OR
new interdomain <domain_id> <address> <len> <dest_domain_id>
```

The parameters for delete, modify, and new are defined as follows:

Parameter	Description
nodeix	The index number of the source node for which a static route is being configured.
address	The source address prefix (19 bytes long).
len	The number of significant bits in the source address prefix.
index	The index number for this static route. This is used in combination with address and len. More than one static route can be configured on a given port with the same address and len values, as long as the index numbers are different.
-port <port>	The number of the port to which the static route is attached.
-vpi <vpi>	The virtual path number to which the static route is attached.
-mtag <mtag>	The metrics tag from the **conf pnni metric** menu that is associated with this static route.
-vpcap (true|false)	**true** means VP capability is set for this PNNI route address. If the signalling interface on which this PNNI route address was created supports VP capability, then the address is advertised by PNNI with VP capability. **false** means VP capability is not set for this PNNI route address. PNNI does not advertise VP capability for this address no matter if the signalling interface on which this PNNI route address was created supports VP capability or not.
domain_id	The domain identifier of this domain (the domain from which the static route is originating).
address	The destination address prefix (19 bytes long).
len	The number of significant bits in the destination address prefix.
dest_domain_id	The domain identifier of the destination domain.

**ATMR Configuration
Commands**

4.4.1.4 Displaying PNNI Exterior Reachability Address Information

This command lets you display information about any statically configured reachability addresses between this domain and other domains. These routes allow two directly connected domains to share reachability information. Enter the following parameters:

```
myswitch::configuration atmroute pnni address> show
PNNI Route Addr Information:
 Node Address                               Plen Index
  1    47.000580ffe1000000f21a0141.000000000000 104  1
       Port VPI Type     Proto Scope VPcap Mtag OperStatus
       1C1  0   exterior mgmt  0     false 0    advertised
ATMR Inter-Domain Route Information:
 Domain Address                             Plen Destn
  1      47.000580ffe1000000f21a0140.002048aabbcc 104  2
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of the node for which the static route is being displayed.
Address	The source address prefix (19 bytes long).
Plen	The number of significant bits in the source address prefix.
Index	The index number for this static route, used in combination with Address and Plen. More than one static route can be configured on a given port with the same Address and Plen values, as long as the index numbers are different.
Port	The number of the port to which the static route is attached.
VPI	The virtual path number to which the static route is attached.
Type	Shows what type of reachability address this is, based on what was configured in the **conf atmr policy** menu. This field is read-only.
Proto	**mgmt** means that the addresses displayed are those addresses that are configured via AMI (configured from SNMP). This field is read-only.
Scope	Shows the scope, which is the highest level at which this reachability information can be advertised. If the address has a scope indicating a level lower than the level of the node, the node will not advertise it. If the address has a scope indicating a level higher than or equal to the level of the node, the node will advertise it to its peer group. The default is 0, which means that the address will be advertised to all levels. This field is read-only.
VPcap	**true** means VP capability is set for this PNNI route address. If the signalling interface on which this PNNI route address was created supports VP capability, then the address is advertised by PNNI with VP capability. **false** means VP capability is not set for this PNNI route address. PNNI does not advertise VP capability for this address no matter if the signalling interface on which this PNNI route address was created supports VP capability or not.

Field	Description
Mtag	The metrics tag from the `conf pnni metric` submenu that is associated with this static route. This field is read-only.
OperStatus	Shows if just the summarized prefix of this address will be announced to the node's peer group, if the entire address will be advertised to the node's peer group, or if this address will not be announced to the node's peer group at all. This field is read-only.
Domain	The domain identifier of this domain (the domain from which the static route is originating).
Address	The destination address prefix (19 bytes long).
Plen	The number of significant bits in the destination address prefix.
Destn	The domain identifier of the destination domain.

You can also display just certain portions of the reachability information. Enter the following parameters:

```
myswitch::configuration atmroute pnni> address show [(pnni | interdomain) [<nodeix>
[<address> [<len> [<index>]]]]]
```

For example, you can show just the PNNI reachability information as follows:

```
myswitch::configuration atmroute pnni address> show pnni
PNNI Route Addr Information:
 Node Address                                   Plen Index
 1    47.000580ffe1000000f21a0141.000000000000 104  1
      Port VPI Type     Proto Scope VPcap Mtag OperStatus
      1C1  0   exterior mgmt  0     false 0    advertised
```

The fields in this display are defined in the same manner as those in the previous example.

You can show just the ATMR inter-domain route information as follows:

```
myswitch::configuration atmroute pnni address> show interdomain
ATMR Inter-Domain Route Information:
 Domain Address                                 Plen Destn
 1      47.000580ffe1000000f21a0140.002048aabbcc 104  2
```

The fields in this display are defined in the same manner as those in the previous example.

If no PNNI addresses have been configured, then the following is displayed:

```
myswitch::configuration atmroute pnni address> show
No route address information is available
No Inter Domain Route information available
```

4.4.2   PNNI Crankback Configuration Commands

These commands let you configure crankback on a switch. During PNNI signalling, a call being processed according to a DTL may encounter a blocked node or link along the designated route. Crankback allows a partial reroute of such a rejected call so that it does not have to be cleared all the way back to the source. Additionally, an indication of the blockage is sent to the originator of the DTL. You can display the list of available subcommands by typing **?** at the **crankback** level.

```
myswitch::configuration atmroute pnni crankback> ?
   show            set
```

4.4.2.1   Displaying the Crankback Setting

This command lets you display the number of times a PNNI call is attempted through crankback on this switch before it is rejected. Enter the following parameters:

```
myswitch::configuration atmroute pnni crankback> show
 Number of tries per call : 2
```

The field in this display is defined as follows:

Parameter	Description
Number of tries per call	The total number of times a PNNI call is attempted and retried through crankback before it is rejected. The default is 2.

4.4.2.2   Configuring the Crankback Setting

This command lets you set the number of times a PNNI call will be retried through crankback on this switch before it is rejected. Enter the following parameters:

```
myswitch::configuration atmroute pnni crankback> set <tries>
```

This parameter is defined as follows:

Parameter	Description
tries	The total number of times a PNNI call is attempted and retried through crankback before it is rejected. The default is 2.

4.4.3 PNNI Interface Configuration Commands

An ATM Forum PNNI node can be bound to a given network-to-network interface (NNI). This is useful when configuring more than one node on a switch. By default, there is one node on a switch configured with the *ForeThought* PNNI protocol and this node is bound to all existing NNIs. The default NNI routing protocol of the default domain dictates the PNNI interface type. You can display the list of available subcommands by typing **?** at the **interface** level.

```
myswitch::configuration atmroute pnni interface> ?
  modify          show
```

4.4.3.1 Modifying a PNNI Interface

This command lets you modify a PNNI interface. Enter the following parameters:

```
myswitch::configuration atmroute pnni interface> modify <port> <vpi>
                [-nodeix <nodeix>]
                [-aggrtoken <aggrtoken>]
                [-cbrw <cbrw>]
                [-rtvbrw <rtvbrw>]
                [-nrtvbrw <nrtvbrw>]
                [-abrw <abrw>]
                [-ubrw <ubrw>]
```

The parameters for modify are defined as follows:

Parameter	Description
port	The port number of the interface to be modified.
vpi	The virtual path number of the interface to be modified.
-nodeix <nodeix>	The index of the node to which the interface is attached.
-aggrtoken <aggrtoken>	The link aggregation token value that is advertised in the Hello protocol. Currently, this parameter is not supported.
-cbrw <cbrw>	The administrative weight of this interface for CBR traffic.
-rtvbrw <rtvbrw>	The administrative weight of this interface for real-time VBR traffic.
-nrtvbrw <nrtvbrw>	The administrative weight of this interface for non real-time VBR traffic.
-abrw <abrw>	The administrative weight of this interface for ABR traffic.
-ubrw <ubrw>	The administrative weight of this interface for UBR traffic.

4.4.3.2 Displaying a PNNI Interface

This command lets you display information about the PNNI interfaces. Enter the following:

```
myswitch::configuration atmroute pnni interface> show
 Port VPI  Node PortID    AggrT VPCap CbrW   RtVbrW  NrtVbrW AbrW    UbrW
 1A1  0    N/A  0x10000000 0     false 5040   5040    5040    5040    5040
 1A2  0    N/A  0x10000001 0     false 5040   5040    5040    5040    5040
 1A3  0    N/A  0x10000002 0     false 5040   5040    5040    5040    5040
 1A4  0    1    0x10000003 0     false 5040   5040    5040    5040    5040
 1A5  0    1    0x10000004 0     false 5040   5040    5040    5040    5040
 1A6  0    1    0x10000005 0     false 5040   5040    5040    5040    5040
 1CTL 0    N/A  0x10000038 0     false 5040   5040    5040    5040    5040
```

The fields in this display are defined as follows:

Field	Description
Port	The port number of the interface.
VPI	The virtual path number of the interface.
Node	The index of the node to which the interface is attached. This shows the configured node index value. The operational node index is displayed under **conf atmr show** and under **conf signalling show atmroute**.
PortID	The internal representation of this port used by ATM Forum PNNI.
AggrT	The link aggregation token value that is advertised in the Hello protocol. This field is read-only and set to 0.
VPCap	`true` means that this interface has VP switching capability and `false` means that it does not. This field is read-only and set to false.
CbrW	The administrative weight of this interface for CBR traffic. The default is 5040.
RtVbrW	The administrative weight of this interface for real-time VBR traffic. The default is 5040.
NrtVbrW	The administrative weight of this interface for non real-time VBR traffic. The default is 5040.
AbrW	The administrative weight of this interface for ABR traffic. The default is 5040.
UbrW	The administrative weight of this interface for UBR traffic. The default is 5040.

You can also display the PNNI interface of a specific port, or a specific port and vpi as follows:

```
myswitch::configuration atmroute pnni> interface show [<port> <vpi>]
myswitch::configuration atmroute pnni> interface show 1a2 0
 Port VPI  Node PortID    AggrT VPCap CbrW   RtVbrW  NrtVbrW AbrW    UbrW
 1A2  0    N/A  0x10000001 0     false 5040   5040    5040    5040    5040
```

The fields in this display are defined in the same manner as those in the previous example.

4.4.4   PNNI Metric Configuration Commands

These commands allow you to configure PNNI metric sets. You can display the list of available subcommands by typing **?** at the **metric** level.

```
myswitch::configuration atmroute pnni metric> ?
  delete          modify          new          show
```

4.4.4.1   Deleting a Metric Set

This command lets you delete a metric set. Enter the following parameters:

```
myswitch::configuration atmroute pnni metric> delete <nodeix> <tag>
(incoming | outgoing) <index>
```

4.4.4.2   Modifying a Metric Set

This command lets you modify a metric set. Enter the following parameters:

```
myswitch::configuration atmroute pnni metric> modify <nodeix> <tag>
(incoming|outgoing) <index>
[cbr] [rtvbr] [nrtvbr] [abr] [ubr] [clpEqual0|clpEqual0Or1]
[-adminw <adminw>] [-pcr <pcr>] [-acr <acr>]
[-ctd <ctd>] [-cdv <cdv>] [-clr <clr>]
[-clr1 <clr1>] [-crm <crm>] [-vf <vf>]
```

4.4.4.3   Creating a Metric Set

This command lets you create a metric set. Enter the following parameters:

```
myswitch::configuration atmroute pnni metric> new <nodeix> <tag>
(incoming|outgoing) <index>
(cbr | rtvbr | nrtvbr | abr | ubr) [clpEqual0|clpEqual0Or1]
[-adminw <adminw>] [-pcr <pcr>] [-acr <acr>]
[-ctd <ctd>] [-cdv <cdv>] [-clr <clr>]
[-clr1 <clr1>] [-crm <crm>] [-vf <vf>]
```

NOTE A class of service (cbr, rtvbr, nrtvbr, abr, or ubr) must be specified when creating a metric set.

The parameters for delete, modify, and new are defined as follows:

Parameter	Description
nodeix	The index number of the node to which this metric set belongs.
tag	Indicates an arbitrary integer that identifies this metric set. 0 is not a valid metric tag. An error is returned if you attempt to use 0.
incoming | outgoing	The direction of the metric set with respect to the owning node.
index	The index number of the metrics being configured within the metric set.
cbr	Indicates that the metric applies to CBR traffic.
rtvbr	Indicates the metric applies to real-time VBR traffic.
nrtvbr	Indicates the metric applies to non real-time VBR traffic.
abr	Indicates that the metric applies to ABR traffic.
ubr	Indicates that the metric applies to UBR traffic.
clpEqual0	Indicates that the advertised GCAC parameters apply to CLP = 0 traffic.
clpEqual0Or1	Indicates that the advertised GCAC parameters apply to CLP = 0+1 traffic.
-adminw <adminw>	The administrative weight value.
-pcr <pcr>	The peak cell rate.
-acr <acr>	The available cell rate.
-ctd <ctd>	The cell transit delay.
-cdv <cdv>	The cell delay variation.
-clr <clr>	The cell loss ratio for CLP=0 traffic.
-clr1 <clr1>	The cell loss ratio for CLP=0+1 traffic.
-crm <crm>	The cell rate margin.
-vf <vf>	The variance factor.

4.4.4.4 Displaying Metric Set Information

This command lets you display metric set information. Enter the following parameters:

```
myswitch::configuration atmroute pnni metric> show
 Node Tag  Direction  Index Cbr   RtVbr NrtVbr Abr   Ubr   ClpType    Adminw
  1   1    incoming   1      false false false  false false clpEqual0  5040
           PCR                ACR                 CTD(usec)    CDV(usec)
           4294967295         4294967295          4294967295   4294967295
           CLR                CLR1                CRM          VF
           4294967295         4294967295          4294967295   4294967295
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of the node to which the metric set belongs.
Tag	The integer that identifies this metric set. 0 means there is no metric associated with the route address or with policy.
Direction	The direction of the metric set with respect to the owning node. Can be either `incoming` or `outgoing`.
Index	The index number of the RAIG being configured within the metric set.
Cbr | RtVbr | NrtVbr | Abr | Ubr	`true` means that the RAIG applies to this type of traffic. `false` means that it does not.
ClpType	`clpEqual0` means that the advertised GCAC parameters apply to CLP=0 traffic. `clpEqual0Or1` means that the advertised GCAC parameters apply to CLP=0+1 traffic.
Adminw	The administrative weight.
PCR	The peak cell rate.
ACR	The available cell rate.
CTD	The cell transit delay.
CDV	The cell delay variation.
CLR	The cell loss ratio for CLP=0 traffic.
CLR1	The cell loss ratio for CLP=0+1 traffic.
CRM	The cell rate margin.
VF	The variance factor.

You can also display information about a specific node, tag, direction, or index as follows:

```
myswitch::configuration atmroute pnni metric> show [<nodeix>] [<tag>]
[(incoming|outgoing)] [<index>]
myswitch::configuration atmroute pnni metric> show 1 1
 Node Tag  Direction  Index Cbr   RtVbr NrtVbr Abr   Ubr   ClpType      Adminw
 1    1    incoming   1     false false false  false false clpEqual0    5040
           PCR             ACR              CTD(usec)      CDV(usec)
           4294967295      4294967295       4294967295     4294967295
           CLR             CLR1             CRM            VF
           4294967295      4294967295       4294967295     4294967295
```

The fields in this display are defined in the same manner as those in the previous example.

If no metrics have been configured, then the following is displayed:

```
myswitch::configuration atmroute pnni metric> show
No metric information is available
```

4.4.5   PNNI Node Configuration Commands

These commands allow you to configure PNNI nodes. You can display the list of available subcommands by typing **?** at the **node** level.

```
myswitch::configuration atmroute pnni node> ?
  admin           delete          new             modify
  show
```

4.4.5.1   Configuring the PNNI Node State

This command lets you bring an existing PNNI node up or take it down. You must administer a node down before you can modify it. Enter the following parameters:

```
myswitch::configuration atmroute pnni node> admin <index> (up | down)
```

4.4.5.2   Deleting a PNNI Node

This command lets you delete an existing PNNI node. Enter the following parameters:

```
myswitch::configuration atmroute pnni node> delete <index>
```

When you delete a node, the switch prompts you with a warning and asks if you really want to delete the node as follows:

```
myswitch::configuration atmroute pnni node> delete 2

Deleting a node will delete all addresses, metrics, profiles, policies
and other information configured for this node. PNNI interfaces currently
attached to this node will be re-attached to any remaining node within
the same domain. If there are no remaining nodes, the interfaces will
display N/A in the 'Node' column until a new node is created in this
domain.

Are you sure you want to delete node [n]? y
```

Entering **n** or pressing <**ENTER**> aborts the command. Entering **y** deletes the node.

4.4.5.3 Creating a PNNI Node

This command lets you create a PNNI node.

NOTE — The maximum number of nodes that can be created per switch is 2. There can be 2 PNNI nodes or one FT-PNNI node and one PNNI node.

Enter the following parameters:

```
myswitch::configuration atmroute pnni node> new <index>
                [-level <level>]
                [-nodeid <nodeid>]
                [-pgid <pgid>]
                [-atmaddr <atmaddr>]
                [-admin (up|down)]
                [-domainname <domainname>]
                [-restrict (true|false)]
                [-ptsehd <ptsehd>]
                [-hellohd <hellohd>] [-hello <hello>]
                [-helloinactf <helloinactf>]
                [-hlinkinact <hlinkinact>]
                [-ptserfr <ptserfr>]
                [-ptselftf <ptselftf>]
                [-rxmt <rxmt>]
                [-avcrpm <avcrpm>] [-avcrmt <avcrmt>]
                [-cdvpm <cdvpm>] [-ctdpm <ctdpm>]
                [-domain <domain>]
                [-forelevel <level>] [-forearea <area>]
```

4.4.5.4   Modifying a PNNI Node

This command lets you modify a PNNI node. Enter the following parameters:

```
myswitch::configuration atmroute pnni node> modify <index>
                [-level <level>]
                [-nodeid <nodeid>]
                [-pgid <pgid>]
                [-atmaddr <atmaddr>]
                [-admin (up|down)]
                [-domainname <domainname>]
                [-restrict (true|false)]
                [-ptsehd <ptsehd>]
                [-hellohd <hellohd>] [-hello <hello>]
                [-helloinactf <helloinactf>]
                [-hlinkinact <hlinkinact>]
                [-ptserfr <ptserfr>]
                [-ptselftf <ptselftf>]
                [-rxmt <rxmt>]
                [-avcrpm <avcrpm>] [-avcrmt <avcrmt>]
                [-cdvpm <cdvpm>] [-ctdpm <ctdpm>]
                [-domain <domain>]
                [-forelevel <level>] [-forearea <area>]
```

> **NOTE**
> If you want to modify **-level**, **-nodeid**, **-pgid**, **-atmaddr**, **-domain**, **-forelevel**, or **-forearea**, you must first administer the node down using **-admin down**.

The parameters for admin, delete, new and modify are defined as follows:

Parameter	Description
index	The index number of the node.
admin	The administrative status of the node. **up** means the node is active. **down** means the node is inactive.
-level <level>	The PNNI hierarchy level of this node. The default is 80.
-nodeid <nodeid>	The ID of this node.
-pgid <pgid>	The peer group ID of the peer group to which this nodes belongs.
-atmaddr <atmaddr>	This node's ATM end system address.

ATMR Configuration Commands

Parameter	Description
-domainname <domainname>	The name of the domain to which this node belongs.
-restrict (true | false)	Indicates whether or not the originating node is restricted only to allow support of SVCs originating or terminating at this node. **true** means that transit capabilities are restricted (i.e., transit connections are not allowed) and **false** means that transit connections are allowed. The default is **false**.
-ptsehd <ptsehd>	The PTSE hold down time, or the minimum interval between updates of any given PTSE. The default is 10 seconds.
-hellohd <hellohd>	The hello hold down time, or the minimum interval between successive hello message transmissions. The default is 10 seconds.
-hello <hello>	The period of time between transmissions of hello messages ("keep alive" messages that two nodes send to one another to verify their existence), in the absence of event-triggered hellos. The default is 15 seconds.
-helloinactf <helloinactf>	The number of hello intervals allowed to pass without receiving a hello message, before a link is declared down. The default is 5 counts.
-hlinkinact <hlinkinact>	The amount of time that a node continues to advertise a horizontal link for which it has not received and processed the logical group node (LGN) horizontal link IG. The default is 120 seconds.
-ptserfr <ptserfr>	The interval between two successive refreshes of a self-originated PTSE in the absence of triggered updates. A node re-originates its PNNI Topology State Elements (PTSEs) at this rate to prevent other nodes from flushing these PTSEs. The default is 1800 seconds.
-ptselftf <ptselftf>	The value used to calculate the initial lifetime of self-originated PTSEs. The initial lifetime is set to the product of the PTSE refresh interval and the PTSE Lifetime Factor. The default is 200%.
-rxmt <rxmt>	The interval at which unacknowledged PTSEs are retransmitted. A PTSE is retransmitted every interval unless explicitly acknowledged through the receipt of either an acknowledgment packet specifying the PTSE instance, or the same instance or a more recent instance of the PTSE by flooding. The default is 5 seconds.
-avcrpm <avcrpm>	The proportional multiplier is expressed as a percentage of the last advertised Available Cell Rate (ACR) on any given PNNI link in the network. If the change in percentage of the ACR on any given PNNI link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant and the topology tables are updated accordingly for that link. This value indicates the threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 50%. The range is 1-99%.
-acvrmt <acvrmt>	The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which the product of the ACR and the proportional multiplier is a very small number. The minimum threshold prevents excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. This value indicates the minimum threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 3%. The range is 1-99%.

Parameter	Description
-domainname <domainname>	The name of the domain to which this node belongs.
-restrict (true | false)	Indicates whether or not the originating node is restricted only to allow support of SVCs originating or terminating at this node. **true** means that transit capabilities are restricted (i.e., transit connections are not allowed) and **false** means that transit connections are allowed. The default is **false**.
-ptsehd <ptsehd>	The PTSE hold down time, or the minimum interval between updates of any given PTSE. The default is 10 seconds.
-hellohd <hellohd>	The hello hold down time, or the minimum interval between successive hello message transmissions. The default is 10 seconds.
-hello <hello>	The period of time between transmissions of hello messages ("keep alive" messages that two nodes send to one another to verify their existence), in the absence of event-triggered hellos. The default is 15 seconds.
-helloinactf <helloinactf>	The number of hello intervals allowed to pass without receiving a hello message, before a link is declared down. The default is 5 counts.
-hlinkinact <hlinkinact>	The amount of time that a node continues to advertise a horizontal link for which it has not received and processed the logical group node (LGN) horizontal link IG. The default is 120 seconds.
-ptserfr <ptserfr>	The interval between two successive refreshes of a self-originated PTSE in the absence of triggered updates. A node re-originates its PNNI Topology State Elements (PTSEs) at this rate to prevent other nodes from flushing these PTSEs. The default is 1800 seconds.
-ptselftf <ptselftf>	The value used to calculate the initial lifetime of self-originated PTSEs. The initial lifetime is set to the product of the PTSE refresh interval and the PTSE Lifetime Factor. The default is 200%.
-rxmt <rxmt>	The interval at which unacknowledged PTSEs are retransmitted. A PTSE is retransmitted every interval unless explicitly acknowledged through the receipt of either an acknowledgment packet specifying the PTSE instance, or the same instance or a more recent instance of the PTSE by flooding. The default is 5 seconds.
-avcrpm <avcrpm>	The proportional multiplier is expressed as a percentage of the last advertised Available Cell Rate (ACR) on any given PNNI link in the network. If the change in percentage of the ACR on any given PNNI link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant and the topology tables are updated accordingly for that link. This value indicates the threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 50%. The range is 1-99%.
-acvrmt <acvrmt>	The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which the product of the ACR and the proportional multiplier is a very small number. The minimum threshold prevents excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. This value indicates the minimum threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 3%. The range is 1-99%.

ATMR Configuration Commands

Parameter	Description
-domainname <domainname>	The name of the domain to which this node belongs.
-restrict (true | false)	Indicates whether or not the originating node is restricted only to allow support of SVCs originating or terminating at this node. **true** means that transit capabilities are restricted (i.e., transit connections are not allowed) and **false** means that transit connections are allowed. The default is **false**.
-ptsehd <ptsehd>	The PTSE hold down time, or the minimum interval between updates of any given PTSE. The default is 10 seconds.
-hellohd <hellohd>	The hello hold down time, or the minimum interval between successive hello message transmissions. The default is 10 seconds.
-hello <hello>	The period of time between transmissions of hello messages ("keep alive" messages that two nodes send to one another to verify their existence), in the absence of event-triggered hellos. The default is 15 seconds.
-helloinactf <helloinactf>	The number of hello intervals allowed to pass without receiving a hello message, before a link is declared down. The default is 5 counts.
-hlinkinact <hlinkinact>	The amount of time that a node continues to advertise a horizontal link for which it has not received and processed the logical group node (LGN) horizontal link IG. The default is 120 seconds.
-ptserfr <ptserfr>	The interval between two successive refreshes of a self-originated PTSE in the absence of triggered updates. A node re-originates its PNNI Topology State Elements (PTSEs) at this rate to prevent other nodes from flushing these PTSEs. The default is 1800 seconds.
-ptselftf <ptselftf>	The value used to calculate the initial lifetime of self-originated PTSEs. The initial lifetime is set to the product of the PTSE refresh interval and the PTSE Lifetime Factor. The default is 200%.
-rxmt <rxmt>	The interval at which unacknowledged PTSEs are retransmitted. A PTSE is retransmitted every interval unless explicitly acknowledged through the receipt of either an acknowledgment packet specifying the PTSE instance, or the same instance or a more recent instance of the PTSE by flooding. The default is 5 seconds.
-avcrpm <avcrpm>	The proportional multiplier is expressed as a percentage of the last advertised Available Cell Rate (ACR) on any given PNNI link in the network. If the change in percentage of the ACR on any given PNNI link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant and the topology tables are updated accordingly for that link. This value indicates the threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 50%. The range is 1-99%.
-acvrmt <acvrmt>	The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which the product of the ACR and the proportional multiplier is a very small number. The minimum threshold prevents excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. This value indicates the minimum threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 3%. The range is 1-99%.

Parameter	Description
-domainname <domainname>	The name of the domain to which this node belongs.
-restrict (true | false)	Indicates whether or not the originating node is restricted only to allow support of SVCs originating or terminating at this node. **true** means that transit capabilities are restricted (i.e., transit connections are not allowed) and **false** means that transit connections are allowed. The default is **false**.
-ptsehd <ptsehd>	The PTSE hold down time, or the minimum interval between updates of any given PTSE. The default is 10 seconds.
-hellohd <hellohd>	The hello hold down time, or the minimum interval between successive hello message transmissions. The default is 10 seconds.
-hello <hello>	The period of time between transmissions of hello messages ("keep alive" messages that two nodes send to one another to verify their existence), in the absence of event-triggered hellos. The default is 15 seconds.
-helloinactf <helloinactf>	The number of hello intervals allowed to pass without receiving a hello message, before a link is declared down. The default is 5 counts.
-hlinkinact <hlinkinact>	The amount of time that a node continues to advertise a horizontal link for which it has not received and processed the logical group node (LGN) horizontal link IG. The default is 120 seconds.
-ptserfr <ptserfr>	The interval between two successive refreshes of a self-originated PTSE in the absence of triggered updates. A node re-originates its PNNI Topology State Elements (PTSEs) at this rate to prevent other nodes from flushing these PTSEs. The default is 1800 seconds.
-ptselftf <ptselftf>	The value used to calculate the initial lifetime of self-originated PTSEs. The initial lifetime is set to the product of the PTSE refresh interval and the PTSE Lifetime Factor. The default is 200%.
-rxmt <rxmt>	The interval at which unacknowledged PTSEs are retransmitted. A PTSE is retransmitted every interval unless explicitly acknowledged through the receipt of either an acknowledgment packet specifying the PTSE instance, or the same instance or a more recent instance of the PTSE by flooding. The default is 5 seconds.
-avcrpm <avcrpm>	The proportional multiplier is expressed as a percentage of the last advertised Available Cell Rate (ACR) on any given PNNI link in the network. If the change in percentage of the ACR on any given PNNI link in the NSAP topology of the network exceeds this percentage threshold, then the change is considered significant and the topology tables are updated accordingly for that link. This value indicates the threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 50%. The range is 1-99%.
-acvrmt <acvrmt>	The minimum threshold is the smallest capacity value that the threshold value for determining the significant change in ACR can take. This minimum value ensures that the threshold value does not become a very small value in cases in which the product of the ACR and the proportional multiplier is a very small number. The minimum threshold prevents excessively frequent NSAP updates resulting from minor changes in ACR when the value of ACR is very low. This value indicates the minimum threshold above which you consider the change in Available Cell Rate (ACR) on any PNNI link to be significant. The default is 3%. The range is 1-99%.

ATMR Configuration Commands

Parameter	Description
-cdvpm <cdvpm>	Indicates the threshold above which you consider the change in Cell Delay Variation (CDV) on any PNNI link to be significant. The default is 25%. The range is 1-99%.
-ctdpm <ctdpm>	Indicates the threshold above which you consider the change in Cell Transit Delay (CTD) on any PNNI link to be significant. The range is 1-99%.
-domain <domain>	The index number of the domain to which this node belongs.
-forelevel <level>	The level of the area used in FORE's hierarchy support. This can be a value between 1 and 127. The default is 5.
-forearea <area>	The ID of the area in FORE's hierarchy support to which this node belongs. This can be a value between 1 and 127. The default is 5.

4.4.5.5 Displaying PNNI Node Information

This command lets you display PNNI node information. Enter the following parameters:

```
myswitch::configuration atmroute pnni node> show
 Node PnniNodeID                                    Level Lowest Rstrn
 1    80:160:47.000580ffe1000000f21a3509.ff1a35090001.00  80    true   false
      PeerGroupID                 PnniAtmAddress
     80:47.000580ffe1000000f2000000 47.000580ffe1000000f21a3509.ff1a35090001.00
      Ptses OperStat AdminStat Shutdown ForeLevel ForeArea DomainId  DomainName
      60    up       up        false    5         5        1
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of this node.
PnniNodeID	The peer group identifier of the peer group to which this node belongs.
Level	The PNNI hierarchy level of the node.
Lowest	`True` means that this node is a lowest level node. This is a read-only field.
Rstrn	Shows whether or not this is a restricted transit node. `true` means that it is and `false` means that it is not.
PeerGroupID	The peer group ID of the peer group to which this node belongs.
PnniAtmAddress	The node's ATM end system address.
Ptses	The number of PTSEs that belong to this node (i.e., the number of PTSEs that are present in this node's database).
OperStat	The current status of this node. `Up` means the node is currently active. `Down` means the node is currently inactive.

Field	Description
AdminStat[1]	Reflects any changes that you have made to the status of the node. Up means you want the node to become active. Down means you want the node to become inactive so you can modify one or more of the parameters.
Shutdown	On a rare occasion, the topology database (TDB) can get into an unrecoverable error state which causes the PNNI logical node to shut down. If this happens, the OperStat field is set to down, the Shutdown field is set to true, and an error message is sent to the console. Then the node quits participating in PNNI activities, such as sending hello packets, performing the database exchange, performing path computation, etc. Other AMI **conf atmr pnni node** commands such as **admin**, **modify**, and **delete** are disabled. Other AMI commands under **display atmr pnni** are disabled (except for **link** and **peer**) and **debug dump atmr pnni tdb** is disabled. To clear the condition and bring the node up again, you must reboot the switch. The default value is false.
ForeLevel	The level of the area to which this PNNI node belongs in the FORE hierarchy. The default is 5.
ForeArea	The ID of the area in the FORE hierarchy to which this PNNI node belongs. The default is 5.
DomainId	The index number of the domain to which this node belongs. The default is 1.
DomainName	The name of the domain to which this node belongs, if one has been assigned.

[1.] When you change the administrative status from down to up, it takes a few seconds for the operational change to occur and to be reflected in the OperStat field. Therefore, it is possible for the display to show the AdminStat as up, but the OperStat as down. If you refresh the display after two or three seconds, the change will have taken place and be reflected.

You can also display timer information, information about significant changes, or information about a specific node. Enter the following parameters:

```
myswitch::configuration atmroute pnni node> show [timer | sigchange] [<nodeix>]
```

To display timer information, enter the following parameters:

```
myswitch::configuration atmroute pnni node> show timer
 Node PtseHD     HelloHD    Hello HelloInactF HLinkInact PtseRfr PtseLftF Rxmt
      (100 msec) (100 msec) (sec)            (sec)      (sec)            (sec)
 1    10         10         15    5          120        1800    200%     5
```

ATMR Configuration Commands

The fields in this display are defined as follows:

Field	Description
Node	The index number of this node.
PtseHD	The minimum interval, in microseconds, between updates of any given PTSE.
HelloHD	The minimum interval, in microseconds, between successive hello message transmissions.
Hello	The period of time, in seconds, between transmissions of hello, or "keep-alive" messages, in the absence of event-triggered hellos.
HelloInactF	The number of hello intervals allowed to pass without receiving a hello message, before a link is declared down.
HLinkInact	The amount of time, in seconds, that a node continues to advertise a horizontal link for which it has not received and processed the logical group node (LGN) horizontal link IG.
PtseRfr	The interval, in seconds, between re-originations of a self-originated PTSE in the absence of triggered updates. A node re-originates its PNNI Topology State Elements (PTSEs) at this rate to prevent other nodes from flushing these PTSEs.
PtseLftF	The value used to calculate the initial lifetime of self-originated PTSEs. The initial lifetime is set to the product of the PTSE refresh interval and the PTSE Lifetime Factor.
Rxmt	The interval, in seconds, at which unacknowledged PTSEs are retransmitted. A PTSE is retransmitted every interval unless explicitly acknowledged through the receipt of either an acknowledgment packet specifying the PTSE instance, or the same instance or a more recent instance of the PTSE by flooding.

To display information about significant changes, enter the following parameters:

```
myswitch::configuration atmroute pnni node> show sigchange
 Node AvcrPm AvcrMt CdvPm  CtdPm
 1    50%    5%     25%    50%
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of this node.
AvcrPm	A change in the Available Cell Rate by this percentage is considered to be significant, provided the change is more than the threshold value specified using AvcrMt.
AvcrMt	The minimum threshold for change in Available Cell Rate after which AvcrPm is used to find the significant change. AvcrMt is expressed as a percentage in Maximum Cell Rate.
CdvPm	A change in Cell Delay Variation by this percentage is considered to be significant.
CtdPm	A change in Maximum Cell Transfer Delay by this percentage is considered to be significant.

> **NOTE** A significant change in any of the above parameters will result in a re-origination of the PTSE.

You can also display information about a specific node. Enter the specific node number as follows:

```
myswitch::configuration atmroute pnni node> show 1
 Node PnniNodeID                                         Level Lowest Rstrn
 1    80:160:47.000580ffe1000000f21a3509.ff1a35090001.00 80    true   false
      PeerGroupID                     PnniAtmAddress
     80:47.000580ffe1000000f2000000 47.000580ffe1000000f21a3509.ff1a35090001.00
      Ptses OperStat AdminStat Shutdown ForeLevel ForeArea DomainId DomainName
      60    up       up        false    5         5        1
```

The fields in this display are defined in the same manner as those listed previously for all of the nodes.

4.4.6   PNNI Profile Configuration Commands

These commands allow you to configure PNNI path computation profiles. You can display the list of available subcommands by typing **?** at the **profile** level.

```
myswitch::configuration atmroute pnni profile> ?
  delete          modify          new             show
```

4.4.6.1 Deleting a PNNI Profile

This command lets you delete an existing PNNI path computation profile. Enter the following parameters:

```
myswitch::configuration atmroute pnni profile> delete <nodeix> <profileix>
```

4.4.6.2 Modifying a PNNI Profile

This command lets you modify an existing PNNI path computation profile. Enter the following parameters:

```
myswitch::configuration atmroute pnni profile> modify <nodeix> <profileix>
               [-minfwdcr <minfwdcr>] [-minrevcr <minrevcr>]
               [-fwdclptype (clpEqual0|clpEqual0Or1)]
               [-revclptype (clpEqual0|clpEqual0Or1)]
               [-fwdclr <fwdclr>] [-revclr <revclr>] [-optctd (true|false)]
               [-optcdv (true|false)] [-loadbal (true|false)]
               [-vponly (true|false)]
```

4.4.6.3 Creating a PNNI Profile

This command lets you create a PNNI path computation profile. Enter the following:

```
myswitch::configuration atmroute pnni profile> new <nodeix> <profileix>/
               (cbr|rtVbr|nrtVbr|abr|ubr)
               [-minfwdcr <minfwdcr>] [-minrevcr <minrevcr>]
               [-fwdclptype (clpEqual0|clpEqual0Or1)]
               [-revclptype (clpEqual0|clpEqual0Or1)]
               [-fwdclr <fwdclr>] [-revclr <revclr>] [-optctd (true|false)]
               [-optcdv (true|false)] [-loadbal (true|false)]
               [-vponly (true|false)]
```

The parameters for delete, new, and modify are defined as follows:

Parameter	Description
nodeix	The index number of the node to which the profile belongs.
profileix	The index number of the profile.
cbr | rtVbr | nrtVbr | abr | ubr	Used only with the **new** command. Indicates for which class of service this profile is used: either cbr, rtVbr, nrtVbr, abr, or ubr.
-minfwdcr \<minfwdcr\>	The minimum forward cell rate requirement.
-minrevcr \<minrevcr\>	The minimum reverse cell rate requirement.
-fwdclptype (clpEqual0 | clpEqual0Or1)	**clpEqual0** means the requested CLR constraint in the forward direction is for the CLP0 stream. **clpEqual0OR1** means the requested CLR constraint in the forward direction is for CLP0+1 stream.
-revclptype (clpEqual0 | clpEqual0Or1)	**clpEqual0** means the requested CLR constraint in the reverse direction is for the CLP0 stream. **clpEqual0OR1** means the requested CLR constraint in the reverse direction is for CLP0+1 stream.
-fwdclr \<fwdclr\>	The forward CLR. Applies only to CBR and VBR class of service.
-revclr \<revclr\>	The reverse CLR. Applies only to CBR and VBR class of service.
-optctd (true | false)	**true** means that the profile is to optimize on CTD and **false** means that it is not. Applies only to CBR and VBR class of service.
-optcdv (true | false)	**true** means that the profile is to optimize on Cell Delay Variation (CDV) and **false** means that it is not. Applies only to CBR and VBR class of service.
-loadbal (true | false)	**true** means that load balancing is activated for the profile and **false** means that it is not. Applies only to ABR and UBR class of service.
-vponly (true | false)	**true** means that the profile is for a VP connection and **false** means that it is not.

ATMR Configuration Commands

4.4.6.4 Displaying PNNI Profile Information

This command lets you display PNNI path computation profile information. See the *ATM Switch Diagnostics and Troubleshooting Manual* for more information about reading this table. Enter the following parameters:

```
myswitch::configuration atmroute pnni profile> show
 Node Profile Service MinFwdCR MinRevCR FwdClpType   RevClpType   FwdCLR  RevCLR
   1    1       cbr     N/A      N/A     clpEqual0    clpEqual0     N/A     N/A
       State     ProfileType LoadBal VPOnly OptCTD OptCDV OptAdmWt NumAvd NumPref
       inactive  mgmtEntry   false   false  false  false  aw        N/A     N/A
 Node Profile Service MinFwdCR MinRevCR FwdClpType   RevClpType   FwdCLR  RevCLR
   1    2       ubr     N/A      N/A     clpEqual0    clpEqual0     N/A     N/A
       State     ProfileType LoadBal VPOnly OptCTD OptCDV OptAdmWt NumAvd NumPref
       computed  cacheEntry  false   false  false  false  aw         0       0
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of the node to which this profile belongs.
Profile	The index number of the profile.
Service	Shows for which class of service this profile is used: either cbr, rtVbr, nrtVbr, abr, or ubr.
MinFwdCR	The minimum forward cell rate requirement.
MinRevCR	The minimum reverse cell rate requirement.
FwdClpType	`clpEqual0` means the requested CLR constraint in the forward direction is for the CLP0 stream. `clpEqual0OR1` means the requested CLR constraint in the forward direction is for CLP0+1 stream.
RevClpType	`clpEqual0` means the requested CLR constraint in the reverse direction is for the CLP0 stream. `clpEqual0OR1` means the requested CLR constraint in the reverse direction is for CLP0+1 stream.
FwdCLR	The forward CLR. Applies only to CBR and VBR class of service.
RevCLR	The reverse CLR. Applies only to CBR and VBR class of service.
State	The current state of the profile.
ProfileType	This is a read-only field. `mgmtEntry` means this profile is dynamically generated. `cacheEntry` means this profile is user-configured.
LoadBal	`true` means that load balancing is activated for the profile and `false` means that it is not. Applies only to CBR and VBR class of service.
VPOnly	`true` means that the profile is for a VP connection and `false` means that it is not.
OptCTD	`true` means that the profile is to optimize on CTD and `false` means that it is not. Applies only to CBR and VBR class of service.

Field	Description
OptCDV	**true** means that the profile is to optimize on Cell Delay Variation (CDV) and **false** means that it is not. Applies only to CBR and VBR class of service.
OptAdmWt	Shows if the profile will optimize on administrative weight.
NumAvd	A read-only field, for cache profiles only, showing the number of avoided links in this profile.
NumPref	A read-only field, for cache profiles only, showing the number of preferred links in this profile.

You can also display information about a specific node or profile as follows:

```
myswitch::configuration atmroute pnni> profile show [<nodeix>] [<profileix>] [advanced]
myswitch::configuration atmroute pnni> profile show 1 1
 Node Profile Service MinFwdCR MinRevCR FwdClpType   RevClpType   FwdCLR  RevCLR
  1    1        cbr     N/A      N/A     clpEqual0    clpEqual0     N/A     N/A
       State      ProfileType LoadBal VPOnly OptCTD OptCDV OptAdmWt NumAvd NumPref
       inactive   mgmtEntry   false   false  false  false  aw       N/A    N/A
```

The fields in this display are defined in the same manner as those in the previous example.

You can also display advanced information about all of the nodes and profiles or a specific node or profile as follows:

```
myswitch::configuration atmroute pnni profile> show advanced 1 2
 Node Profile Service MinFwdCR MinRevCR FwdClpType   RevClpType   FwdCLR  RevCLR
  1    2        ubr     N/A      N/A     clpEqual0    clpEqual0     N/A     N/A
       State      ProfileType LoadBal VPOnly OptCTD OptCDV OptAdmWt NumAvd NumPref
       computed   cacheEntry  false   false  false  false  aw       0      0

       No avoided link information for this profile
       No preferred link information for this profile
```

The fields in this display are defined in the same manner as those in the previous example. There is additional information given about avoided or preferred links, if any such information is available.

If no PNNI profiles have been configured, then the following is displayed:

```
      myswitch::configuration atmroute pnni profile> show
      No profile information is available
```

4.5   PNNI Policy Prefix Configuration Commands

These commands allow you to configure a policy prefix for a node. A policy allows you to impose rules on how a node propagates information to its peer group about other peer groups. Each policy consists of an ATM address prefix, one of three actions associated with the address, and an optional metric tag. (For more information about policy, see the *ATM Switch Network Configuration Manual* for your switch.) You can display the list of available subcommands by typing **?** at the **policy** level.

```
myswitch::configuration atmroute policy> ?
  delete          modify          new             show
```

4.5.1   Deleting a PNNI Policy Prefix

This command allows you to delete an existing policy prefix. Enter the following parameters:

```
myswitch::configuration atmroute policy> delete <nodeix> <address> <len>
```

4.5.2   Modifying a PNNI Policy Prefix

This command allows you to modify an existing policy prefix. Enter the following parameters:

```
myswitch::configuration atmroute policy> modify <nodeix> <address> <len>
        [-type (internal | exterior | both)]
        [-action (summary | suppress | advertise)]
        [-mtag <tag>]
```

4.5.3   Creating a PNNI Policy Prefix

This command allows you to create a new policy prefix. Enter the following parameters:

```
myswitch::configuration atmroute policy> new <nodeix> <address> <len>
        [-type (internal | exterior | both)]
        [-action (summary | suppress | advertise)]
        [-mtag <tag>]
```

The parameters for delete, modify, and new are defined as follows:

Parameter	Description
nodeix	The index number of the node to which the policy applies.
address	The 19-byte address prefix of the policy.
len	The address prefix length of the policy.
-type (internal | exterior | both)	The type of reachable addresses to which the policy applies. `internal` means the policy applies to internal reachable addresses within that node's FORE hierarchy area. `exterior` means the policy applies to exterior reachable addresses outside of that node's FORE hierarchy area. `both` means the policy applies to both internal and exterior addresses. The default is `both`.
-action (summary | suppress | advertise)	The policy for exchanging dynamic reachability address prefixes. Addresses matching a `summary` policy cause just the summarized prefix of the address to be announced to the node's peer group. Addresses matching a `suppress` policy are <u>not</u> announced to the node's peer group at all. Addresses matching an `advertise` policy cause the entire address to be announced to the node's peer group. An `advertise` policy is only used to advertise exceptions to a `suppress` policy. The default is `summary`.
-mtag <tag>	The metric tag associated with this policy. Look in the `Tag` field under **conf atmroute pnni metric show** to find the appropriate tag number. The default is 0.

4.5.4   Displaying PNNI Policy Prefix Information

This command allows you to display information about existing policy prefixes. Enter the following parameters:

```
myswitch::configuration atmroute policy> show
 Node Address                                  Plen Type     Action
 1    47.000580ffe1000000f21a355200204806790b00 72   both     summary
      State      MTag
      inactive   0
```

The fields in this display are defined as follows:

Field	Description
Node	The index number of the node to which this policy applies.
Address	The address prefix of the policy.
Plen	The address prefix length of the policy (the significant part of the address).
Type	The type of reachable addresses to which the policy applies. `internal` means the policy applies to internal reachable addresses within that node's FORE hierarchy area. `exterior` means the policy applies to exterior reachable addresses outside of that node's FORE hierarchy area. `both` means the policy applies to both internal and exterior addresses.
Action	The action to take when using this policy. Addresses matching a `summary` policy cause just the summarized prefix of the address to be announced to the node's peer group. Addresses matching a `suppress` policy are <u>not</u> announced to the node's peer group at all. Addresses matching an `advertise` policy cause the entire address to be announced to the node's peer group.
State	Indicates the current state of this policy. If there are no reachable addresses that match this policy, then the state is listed as `inactive`. When a summary policy is matched and is being used, the state is listed as `summarizing`. When a suppress policy is matched and is being used, the state is listed as `suppressing`. When an advertise policy is matched and is being used, the state is listed as `advertising`. When a policy is being superseded by another policy, the state is listed as `N/A`.
Mtag	The metric tag associated with this policy. See the `Tag` field under **conf atmroute pnni metric show**.

You can also display specific policy information as follows:

```
myswitch::configuration atmroute policy> show [<nodeix> [<address> <len>]]
myswitch::configuration atmroute policy> show 1
 Node Address                                  Plen Type     Action
 1    47.000580ffe1000000f21a355200204806790b00 72   both     summary
      State      MTag
      inactive   0
```

The fields in this display are defined in the same manner as those in the previous example.

4.6   SPANS ATM Route Configuration Commands

These commands allow you to modify various aspects of SPANS-NNI on a switch. You can display the list of available subcommands by typing **?** at the **spans** level.

```
myswitch::configuration atmroute spans> ?
  area              border          show
```

> **NOTE**
>
> In order for part of a FORE ATM cloud to be a hierarchical SPANS area, all switches in that cloud must be running a software version that is at least *ForeThought* 4.0 or greater. Otherwise, SPANS connectivity between the hierarchical area and the SPANS area will be lost.

4.6.1   Setting the SPANS Area ID

This command enables you to set the SPANS area ID. Enter the following parameters:

```
myswitch::configuration atmroute spans> area <area>
```

This parameter is defined as follows:

Parameter	Description
area[1]	The ID of the area in the SPANS routing hierarchy to which this switch belongs. This number goes into the most significant byte of the SPANS ATM address. The default area ID for all switches is 242 in decimal.

[1]. The switch must be rebooted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

4.6.2   Setting the SPANS-NNI Border Switch Functionality

A switch that has a link to another switch belonging to a different SPANS area is considered a border switch. A border switch advertises reachability to its area to switches outside of its area, but it does not share its area's topology with the other switches. You should enable border switch functionality on all switches that are on the outside edges of all of the areas that you have established. This command lets you designate whether or not this switch will act as a SPANS-NNI border switch. Enter the following parameters:

```
myswitch::configuration atmroute spans> border (enable | disable)
```

This parameter is defined as follows:

Parameter	Description
enable | disable[1]	Entering `enable` (and rebooting) means that this switch will be a SPANS border switch. Entering `disable` (and rebooting) means that this switch will not be a SPANS border switch. The default is disable.

[1.] The switch software must be restarted for this command to take effect. Therefore, you must be in a local AMI session to perform this command.

4.6.3   Displaying SPANS-NNI Parameters

This command let you display all of the SPANS-NNI topology parameters. Enter the following parameters:

```
myswitch::configuration atmroute spans> show
SPANS Area ID          242
SPANS NNI border switch functionality is disabled
```

The fields in this display are defined as follows:

Field	Description
SPANS Area ID	The ID of the area in the SPANS routing hierarchy to which this switch belongs. This number goes into the most significant byte of the SPANS ATM address.
SPANS NNI border switch functionality is disabled	If this functionality is `enabled`, this switch is a SPANS-NNI border switch. If it is `disabled`, this switch is not a SPANS-NNI border switch.

CHAPTER 5   Board Configuration Commands

This submenu lets you configure default values for the switch board. Type **?** at the **board** level to display the available subcommands.

```
myswitch::configuration board> ?
  show
```

The following is displayed on an ASX-4000 and on an LE 25:

```
myswitch::configuration board> ?
  mcastspaces          setmcastspace        show
```

5.1   Displaying the Multicast Space Models

This command allows you to display the different types of multicast space models to be used in conjunction with the memory models for Series LE network modules or for the logical network modules on Series 1 port cards. The following models are available on an LE 25:

```
myswitch::configuration board> mcastspaces
MulticastIndex    Name
          1       mcast256
          2       mcast512
```

The following models are available on an ASX-4000:

```
myswitch::configuration board> mcastspaces
MulticastIndex    Name
          1       mcast512
          2       mcast1K
          3       mcast2K
          4       mcast8K
```

The fields in this display are defined as follows:

Field	Description
MulticastIndex	The multicast space model index number. This is the number to enter for the *<index>* parameter when using the **conf board setmcastspace** command.
Name	Indicates the amount of multicast space for this model.

5.2   Configuring the Multicast Space

This command lets you configure the amount of multicast space on a switch board (fabric) for an ASX-4000 or an LE 25 only. This command dictates which memory model (under **conf module traffic pc1 models**) can be used on the logical network modules on the OC-12c Series 1 port cards that are installed in that board, and which memory model (under **conf module traffic le models**) can be used on the LE 25 network modules. If you change the multicast space, the memory models are changed to a compatible model (i.e., a model with the name that matches the mcastspace size) when the switch software is restarted. Enter the following parameters:

```
myswitch::configuration board> setmcastspace <board> <index>
```

These parameters are defined as follows:

Parameter	Description
board	The number of the switch board on an ASX-4000 or LE 25 on which the amount of multicast space is being configured. On an ASX-4000, can be 1, 2, 3, or 4, with 1 being the leftmost of the switch board slots. On an LE 25, this value is always 1.
index	The index number from **conf board mcastspaces** which shows the amount of multicast space that you want to use for this ASX-4000 or LE 25 switch board. The **mcast8K** option is only available on boards that contain <u>only</u> Series 1 OC-48c port cards. The default is index 1 (mcast512) for all Series 1 port cards. The default for an LE 25 is index 1 (mcast256).

If you change the amount of multicast space, the switch warns you about the effects of the change and prompts you to restart the software as follows:

```
myswitch::configuration board> setmcastspace 1 2
   WARNING: For correct operation, all OC3/OC12 Series PC1
   port cards on board 1 must use memory models
   that support mcast1K. OC3/OC12 Series PC1 port
   cards may be autoconfigured to a new memory model
   that supports this multicast space.

   Changing the multicast space model on a board will take
   effect after restarting the switch software.

   Set multicast space index to 2 [y]? y
   Restart the switch software [n]? y
```

5.2.1   Displaying the Board Configuration

This command shows the current configuration of the switch board (switch fabric). Enter the
following parameters:

```
myswitch::configuration board> show
Board  Model     Rev  S/N       NMs  Multicast
4      asx1000   J    8014      2    N/A
```

The following is displayed on an ASX-4000:

```
myswitch::configuration board> show
Board  Model     Rev  S/N       NMs  Multicast
1      asx4000   5    98200044  4    mcast512
2      asx4000   5    98200040  2    mcast512
```

The following is displayed on an LE 25:

```
myswitch::configuration board> show
Board  Model     Rev  S/N       NMs  Multicast
1      le25      0    132       4    mcast256
```

The fields in this display are defined as follows：

Field	Description
Board	The number of the slot in which the board (switch fabric) is installed.
Model	The type of switch this is.
Rev	The manufacturing revision number.
S/N	The serial number of this switch board.
NMs	The number of network modules installed in this board.
Multicast	The multicast space configured for this switch board. This field only applies to an ASX-4000 or an LE 25.

You can also display information about a particular switch board in an ASX-4000 as follows:

```
myswitch::configuration board> show [<board>]
localhost::configuration board> show 2
Board  Model     Rev  S/N       NMs  Multicast
2      asx4000   5    98200040  2    mcast512
```

The fields in this display are defined in the same manner as those in the previous example.

CHAPTER 6   CEC Configuration Commands

The **cec** commands let you configure the Timing Control Module (TCM) on a CEC-Plus. The CEC-Plus is an environmental/timing management subsystem. To display the **cec** commands, a TCM must be installed in the switch. You can display the list of available subcommands by typing **?** at the **cec** level.

```
myswitch::configuration cec> ?
  alarms>         slotx>          sloty>          timing>
```

6.1   Alarms Configuration Commands

This submenu lets you configure alarm conditions that are a result of various environmental and synchronization timing conditions. You can display the list of available subcommands by typing **alarms ?** at the **cec** level.

```
myswitch::configuration cec> alarms ?
  disable         enable          relays>         show
```

6.1.1   Disabling an Alarm

This command lets you disable an alarm. Enter the following parameters:

```
myswitch::configuration cec alarms> disable (major | minor) <alarm type>
```

These parameters are defined as follows:

Parameter	Description
major | minor	Designates whether the alarm type causes a major alarm or a minor alarm when that condition occurs.
alarm type	Indicates the alarm condition. Valid parameters are displayed in the `AlarmType` field when the command string `conf alarms show` is entered.

For example, to disable an overtemperature condition that is detected by the overtemperature sensor as a minor alarm, enter the following parameters:

```
myswitch::configuration cec alarms> disable minor tempSensorOverTemp
```

To verify that the change has taken effect, you can display the alarms:

```
myswitch::configuration alarms> show
 AlarmType                     AlarmStatus   MinorAlarm   MajorAlarm
 powerSupplyInputFailed        active        disabled     enabled
 powerSupplyOutputFailed       active        disabled     enabled
 fanBankFailed                 active        disabled     enabled
 tempSensorOverTemp            inactive      disabled     enabled
 Major alarm relay status: on
 Minor alarm relay status: off
```

6.1.2   Enabling an Alarm

This command lets you enable an alarm. Enter the following parameters:

```
myswitch::configuration cec alarms> enable (major | minor) <alarm type>
```

These parameters are defined as follows:

Parameter	Description
major | minor	Designates whether the alarm type causes a major alarm or a minor alarm when that condition occurs.
alarm type	Indicates the kind of alarm condition. Valid parameters are displayed in the AlarmType field when the command string **conf alarms show** is entered at the prompt.

For example, to enable an overtemperature condition that is detected by the overtemperature sensor as a major alarm, enter the following parameters:

```
myswitch::configuration cec alarms> enable major tempSensorOverTemp
```

To verify that the change has taken effect, you can display the alarms:

```
myswitch::configuration alarms> show
 AlarmType                    AlarmStatus   MinorAlarm   MajorAlarm
 powerSupplyInputFailed       active        disabled     enabled
 powerSupplyOutputFailed      active        disabled     enabled
 fanBankFailed                active        disabled     enabled
 tempSensorOverTemp           inactive      disabled     enabled
 Major alarm relay status: on
 Minor alarm relay status: off
```

6.1.3   Configuring an Alarm Relay

These commands let you configure or display alarm relays. You can show the list of available subcommands by typing **relays ?** at the **alarms** level. Enter the following parameters:

```
myswitch::configuration cec alarms> relays ?
   set              show
```

6.1.3.1 Setting an Alarm Relay

This command lets you set an alarm relay. Enter the following parameters:

```
myswitch::configuration cec alarms relays> set <relay number> (major | minor | unused)
```

These parameters are defined as follows:

Parameter	Description
relay number	Indicates the number which corresponds to one of the five alarm relay LEDs on the front of the SCP.
major|minor | unused	Indicates if you want the alarm relay to be major, minor, or unused.

6.1.3.2 Displaying the Alarm Relays

This command lets you display the alarm relays. Enter the following parameters:

```
myswitch::configuration cec alarms relays> show
 Relay            Alarm           Current
 Number           Function         State
    1             major             on
    2             minor             on
    3             unused            off
    4             unused            off
    5             major             on
```

The fields in this display are defined are defined as follows:

Field	Description
relay number	Shows the number which corresponds to one of the five alarm relay LEDs on the front of the SCP.
Alarm Function	Shows whether the alarm relay has been configured as major, minor, or unused.
Current State	Shows whether the alarm relay is on (in a state of alarm) or off (not in a state of alarm).

6.1.4   Displaying Alarm Conditions

This command lets you display the status of all alarms. Enter the following:

```
myswitch::configuration cec alarms> show
 AlarmType                   AlarmStatus    MinorAlarm    MajorAlarm
 powerSupplyInputFailed      active         disabled      enabled
 powerSupplyOutputFailed     active         disabled      enabled
 fanBankFailed               active         disabled      enabled
 tempSensorOverTemp          inactive       disabled      enabled
 Major alarm relay status: on
 Minor alarm relay status: off
```

The fields in this display are defined as follows:

Field	Description
AlarmType	Displays the name of the alarm.
AlarmStatus	Shows whether the state of the alarm is active (alarming) or inactive (not alarming). An alarm is active if the underlying condition is detected. For power supplies, the input failed alarm condition is active if the input voltage is not within the nominal range for the supply. This does not necessarily mean that an output failure will result. A power supply output failure condition is active if any power supply is failing or if it is physically removed.
MinorAlarm	`disabled` means that this alarm type will not cause a minor alarm. `enabled` means that this alarm type causes a minor alarm.
MajorAlarm	`disabled` means that this alarm type will not cause a major alarm. `enabled` means that this alarm type causes a major alarm.
Major Alarm relay status	`off` means that no major alarms are currently active. `on` means that one or more major alarms are currently active. Look at the `AlarmStatus` field to see which condition is in a state of alarm.
Minor Alarm relay status	`off` means that no minor alarms are currently active. `on` means that one or more minor alarms are currently active. Look at the `AlarmStatus` field to see which condition is in a state of alarm.

6.2 TCM Configuration Commands

The **slotx** and **sloty** submenus allow you to choose a TCM for further configuration. To configure parameters on the TCM in slot X (the top slot), type **slotx** at the **cec** level. To configure the TCM in slot Y (the bottom slot), type **sloty** at the **cec** level.

Once an individual TCM has been selected for configuration, you can configure IP parameters, change the name of the TCM, display information about the TCM, and so on. You can display the list of available subcommands by typing **?** at the **slotx** or **sloty** sublevel. The commands for both levels are the same, but they are only described once with examples that read **slotx**.

```
myswitch::configuration cec slotx> ?
  ip>             name            rs232>          show
  snmp>           timezone
```

6.2.1 IP Configuration Commands

> **NOTE**
>
> The ie0 and ie1 interfaces of all installed TCMs and the ie0 interface of all installed SCPs must all reside on the same subnet.

These commands let you change the IP configuration of the TCM's interfaces. You can display the list of available subcommands by typing **ip ?** at the **configuration** level.

```
myswitch::configuration cec slotx ip> ?
  address         admin           broadcast       route>
  show
```

6.2.1.1 Configuring the IP Address

This command lets you configure an IP address, netmask, and administrative status for one of the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip> address <interface> <address> [<netmask>]
[(up)]
```

These parameters are defined as follows:

Parameter	Description
interface	Indicates the name of the IP interface to be managed.
address	Indicates the IP address for this interface.
netmask	The subnet mask for this IP interface. It should be entered in dotted decimal notation (e.g., 255.255.255.0). If you accidentally enter a duplicate address for the same subnet as an existing one, the switch rejects the duplicate and prints an error message. If you do not provide a netmask, then a default netmask is used.
up	Entering up brings the designated interface on-line as soon as you enter this command.

6.2.1.2 Configuring the IP State

This command lets you enable or disable the IP interfaces on the TCM. Enter the following parameters:

```
myswitch::configuration cec slotx ip> admin <interface> (up|down)
```

These parameters are defined as follows:

Parameter	Description
interface	Indicates the name of the IP interface to be managed.
up | down	Entering up enables the designated interface. Entering down disables the designated interface.

6.2.1.3 Configuring the IP Broadcast Address

This command allows you to modify the broadcast address for one of the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip> broadcast <interface> (0|1)
```

These parameters are defined as follows:

Parameter	Description
interface	Indicates the name of the IP interface.
0|1	Indicates the IP broadcast type for this interface. This is the host portion of the IP address that is used for routing. Entering 1 causes the host portion of the IP address to be set to all 1s. Entering 0 causes the host portion of the IP address to be set to all 0s.

6.2.1.4 Configuring IP Routes

This command allows you to add a static IP route to the local IP routing table, delete a static IP route from the local IP routing table, or list the current static IP routes in the local IP routing table for one of the TCM's IP interfaces. You can display the list of available subcommands by typing **route ?** at the **ip** level.

```
myswitch::configuration cec slotx ip> route ?
   new              delete           show
```

6.2.1.4.1   Adding an IP Route

This command lets you create an IP route for one of the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip route> new (default|<destination-ipaddress>)
<gateway> [<metric>] [(host | net)]
```

These parameters are defined as follows:

Parameter	Description
default	This parameter must be specified to create a default route.
destination-ipaddress	Indicates the destination IP network number.
gateway	Indicates the gateway address to the destination IP network number.
metric	Indicates the number of hops to the destination IP network. The default value of 1 is used if no value is entered. If 1 is specified, the route is created with the RTF_GATEWAY flag.
host|net	Using host indicates this is a host-specific route with the destination being a specific node's IP address. Using net indicates this is a network-specific route with the destination being a network IP address. The default value of net is used if no value is entered.

6.2.1.4.2   Deleting an IP Route

This command lets you delete an IP route from one of the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip route> delete (default|<destination-ipaddress>)
<gateway>
```

These parameters are defined as follows:

Parameter	Description
default	A default must be specified to delete a default route.
destination-ipaddress	Indicates the destination IP network number.
gateway	Indicates the gateway address to the destination IP network number.

6.2.1.4.3   Showing the IP Routes

This command lets you display the current IP routes for the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip route> show
Destination      Gateway          Metric        Interface       Flags
default          198.29.31.75     1             ie0             G
127.0.0.1        127.0.0.1        0             lo0
169.144.85.3     198.29.31.75     1             ie0             G
198.29.31.0      198.29.31.28     0             ie0
```

The fields in this display are defined as follows:

Field	Description
Destination	Indicates the destination IP network number.
Gateway	Indicates the gateway address to the destination IP network number.
Metric	Shows the number of hops to the destination IP network. The default is 1
Interface	Shows the local IP interface used to get to the destination IP network.
Flags	Shows H if the route is host-specific (created with the RTF_HOST flag set). Shows G if the route is network-specific (created with the RTF_GATEWAY flag set).

6.2.1.5 Displaying the IP Interface Configuration

This command allows you to display information about the configuration of the TCM's IP interfaces. Enter the following parameters:

```
myswitch::configuration cec slotx ip> show
interface  state       address        netmask        broadcast
ie0        up          169.144.28.125 255.255.255.0  169.144.28.255
ie1        up          169.144.28.225 255.255.255.0  169.144.28.255
lo0        up          127.0.0.1      255.0.0.0      N/A
```

The fields in this display are defined as follows:

Field	Description
interface	Indicates the name of the IP interface.
state	Lists the administrative state of the IP interface.
address	Displays the IP address of the IP interface.
netmask	Shows the netmask address of the IP interface.
broadcast	Indicates the broadcast address of the IP interface.

You may also designate a single interface to be displayed by entering **show** and the specific interface name at the prompt as follows:

```
myswitch::configuration cec slotx ip> show ie0
interface  state       address        netmask        broadcast
ie0        up          169.144.28.125 255.255.255.0  169.144.28.255
```

The fields in this display are defined in the same manner as those listed above in the example for showing the configuration of all of the IP interfaces.

6.2.2 Setting or Changing the TCM's Name

This command lets you set or change the name of the TCM. Enter the following parameters:

```
myswitch::configuration cec slotx> name <name>
```

This parameter is defined as follows:

Parameter	Description
name	Indicates the new system name for the TCM.

For example, to set the TCM's name to linus, enter the following parameters:

```
myswitch::configuration cec slotx> name linus
```

6.2.3   Serial Port Configuration

This command lets you display configuration information for the RS-232 serial port on the front panel of the TCM. You can display the available subcommand by typing **?** at the **rs232** level.

```
myswitch::configuration cec slotx rs232> ?
   show
```

6.2.3.1   Displaying Serial Port Information

This command allows you to display the settings for the RS-232 serial port on the TCM. Enter the following parameters:

```
myswitch::configuration cec slotx rs232> show
Port    Type       Speed  Flow    Bits  Stops    Parity
A       rs232       9600   none      8   one      none
```

The fields in this display are defined as follows:

Field	Description
Port	Shows the physical port designation.
Type	Shows the signalling standard used.
Speed	Shows the receive/transmit rate in bits per second.
Flow	Shows the type of flow control implemented on the given port.
Bits	Shows the number of bit times in a single character.
Stops	Shows the number of stop bits in a character frame.
Parity	Shows the parity setting for the ports.

If no TCM is installed, then the following is displayed:

```
myswitch::configuration cec slotx rs232> show
No TCM found in this slot
```

6.2.4 Displaying TCM Information

This command lets you display information about the TCM to which you are logged in. Enter the following parameters:

```
myswitch::configuration cec slotx> show
TCM 'linus', Type cec-plus, up 0 days 18:15
Hardware version A, Software version S_ForeThought_5.2.0 (1.14891)
Serial number                      30
Slot                               X
State                              active
Time zone
External Inputs (1-5)              off off off off off
Change to active operation occurred at Apr 16  17:16:56 1998
ESI module is present.
```

The fields in this display are defined as follows:

Field	Description
TCM	This TCM's system name.
Type	Shows what type of TCM this is.
up	The amount of time (in days, hours, and minutes) since this TCM has been rebooted.
Hardware version	The hardware version of this TCM.
Software Version	The software version being used by this TCM.
Serial number	The serial number of this TCM card.
Slot	The slot in which this TCM resides. X indicates the top slot and Y indicates the bottom slot.
State	The current state of this TCM. active means this TCM is the controller and is functioning properly. standby means this is the standby TCM. offline means the timing has failed on this TCM.
Time zone	The time zone configured for this TCM. If this field reads N/A, then this value has not been configured yet.
External Inputs	The current state of each of the five front panel external inputs from the viewpoint of this TCM. Can be either on (input being received) or off (no input).
active | standby | offline	Shows when this TCM control software switched to the current operating mode.
ESI module	The External Synchronization Interface (ESI) card status from the viewpoint of this TCM. This should always be present. absent indicates a failure of the ESI interface logic.
Other TCM status	If another TCM is not plugged in, nothing is displayed. If another TCM is plugged in, shows the software status of the other TCM from the viewpoint of this TCM. If this TCM is the controller, normal means the standby TCM is actively updating and waiting to be called into service, and unknown means the standby TCM is down. If this TCM is the standby, active means the controller TCM is functioning normally.

6.2.5   SNMP Configuration Commands

These commands enable you to manage the SNMP communities and traps. You can display
the list of available subcommands by typing **snmp ?** at the **configuration** level.

```
myswitch::configuration cec slotx> snmp ?
trap>
```

6.2.5.1 Configuring SNMP Traps

These commands help you to manage SNMP traps. You can display the list of available sub-
commands by typing **trap ?** at the **snmp** level.

```
myswitch::configuration cec slotx snmp> trap ?
delete          new             show
```

6.2.5.1.1   Deleting an SNMP Trap Entry

This command allows you to delete an existing SNMP trap destination. Before deleting a trap
that may need to be recreated later, show the list of current SNMP traps and either copy and
save the screen or write down the trap destinations. You will also need to show the list of cur-
rent SNMP traps in order to find the number of the trap to be deleted. Enter the following
parameters to delete a trap entry:

```
myswitch::configuration cec slotx snmp trap> delete <trap>
```

The parameter for deleting is defined as follows:

Parameter	Description
trap	Indicates the number of the trap destination in the list of current SNMP traps that is to be removed.

For example, to delete trap 198.29.31.130, first list the traps to find its number and copy the
address in case you want to recreate it later:

```
myswitch::configuration cec slotx snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
5       198.29.31.130
```

Then enter the following parameters:

```
myswitch::configuration cec slotx snmp trap> delete 5
```

You can display the list again to verify that the trap has been deleted:

```
myswitch::configuration cec slotx snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
```

6.2.5.1.2   Creating an SNMP Trap Entry

This command allows you to specify a host to which an TCM can send SNMP traps. Enter the IP address of the SNMP trap destination to be added. Repeat this for as many SNMP trap destinations as needed. Traps are active as soon as they are set. Enter the following parameters:

```
myswitch::configuration cec slotx snmp trap> new <ipaddress>
```

The parameter for specifying is defined as follows:

Parameter	Description
ipaddress	Indicates the IP address of the trap destination to be created.

6.2.5.1.3   Displaying the SNMP Trap Entries

This command enables you to list all of the current SNMP traps. Enter the following:

```
myswitch::configuration cec slotx snmp trap> show
Trap    Destination
1       192.88.243.18
2       198.29.16.14
3       198.29.16.18
4       198.29.23.39
5       198.29.31.130
```

If no SNMP traps have been configured, the following message is displayed:

```
No trap information is available
```

6.2.6   Setting or Changing the Timezone

This command enables you to set or change the timezone on the TCM. Enter the following:

```
myswitch::configuration cec slotx> timezone <timezone>
```

The parameter for setting or changing is defined as follows:

Parameter	Description
timezone	The time zone configured for this TCM. The TCM supports and automatically converts from Standard to Daylight Savings time for the following time zones: EST5EDT (Eastern Standard Time), CST6CDT (Central Standard Time), MST7MDT (Mountain Standard Time), PST8PDT (Pacific Standard Time), AKST9AKDT (Alaska Standard Time).

Locations outside of the time zones listed above must supply the following POSIX standard 1003.1-1988 formula for switching between Daylight Savings Time and Standard Time:

stdoffset [dst[offset][, start[/time], end[/time]]]

The time zones have the following meanings:

Field	Description
std and dst	Indicates 3 or more bytes that designate standard (std) or Daylight Savings Time (dst). Only std is required; if dst is omitted, then it does not apply in this location. Can use uppercase or lowercase letters and any characters, except a leading colon(:), digits, comma (,), minus (-), plus (+), and ACSII NUL.
offset	Indicates the value to add to local time to equal Greenwich Mean Time. offset is of the form: **hh[:mm[:ss]]** Hour (hh) is required and can be a single digit between 0 and 24. Minutes (mm) and seconds (ss) are optional and are between 0 and 59. If no offset follows dst, it is assumed to be one hour ahead of std. If preceded by a "−", the time zone is east of the Prime Meridian; otherwise it is west (with an optional "+")
start[/time], end[/time]	start indicates the date when the change occurs from std to dst. end indicates the date when you change back. Both start and end are of the form: **Mm.n.d** d is the d-th day $(0 \le d \le 6)$ of week n of month m of the year $(1 \le n \le 5, 1 \le m \le 12)$, where week 5 is the last d day in month m, which can occur in either the fourth or the fifth week). Week 1 is the first week in which the d-th day occurs. Day 0 is Sunday. time is of the same format as offset, except that no leading "−" or "+" is allowed. If time is not entered, the default of 02:00:00 is used.

6.3   CEC Timing Configuration

These commands let you configure external synchronization timing. Type **timing ?** at the **cec** level to display the list of available subcommands.

```
myswitch::configuration cec> timing ?
  admin         bits>         failover>       mode
  primary       references    refqual         revertive>
  secondary     show
```

6.3.1   Configuring the Timing Reference's Administrative Status

This command lets you configure the administrative status of the timing reference that is provided to the ESI. You can get to this level by entering **admin** at the **timing** level. Enter the following parameters:

```
myswitch::configuration timing> admin fabric (1 | 2 | 3 | 4)
(primary | secondary) (up | down)

or

admin (bits1 | bits2) (up | down)
```

These parameters are defined as follows:

Parameter	Description
fabric (1 | 2 | 3 | 4)	Indicates the switch fabric (1 - first (leftmost slot), 2 - second slot, 3 - third slot, 4 fourth slot) of the timing source for which you are changing the administrative status.
(primary | secondary)	Indicates for which timing source you are changing the administrative status.
(bits1 | bits2)	Indicates for which BITS timing source you are changing the administrative status.
(up | down)	up means you are enabling this timing source. **down** means you are disabling this timing source.

6.3.2   BITS Timing Configuration Commands

These commands let you configure the timing input and output for the BITS clock. You can display the list of available subcommands by entering **bits ?** at the **timing** level. Enter the following:

```
myswitch::configuration cec timing> bits ?
   coding              framing             level
```

6.3.2.1   BITS Coding Configuration

This command lets you configure the coding of the DS1 BITS interface or the E1 BITS interface for this TCM. You can get to this level by entering **coding** at the **bits** level. Enter the following parameters:

```
myswitch::configuration cec timing bits> coding (ami | b8zs) for DS1 interface
                            Note: only hdb3 supported for E1 interface
```

These parameters are defined as follows:

Parameter	Description
ami	Indicates that Alternate Mark Inversion (AMI) coding should be used for the BITS interface on the TCM. This means zeros are represented by 01 during each bit cell, and ones are represented by 11 or 00, alternately, during each bit cell. This technique requires that the sending device maintain ones density. Ones density is not maintained independent of the data stream.
b8zs	Indicates that Binary 8-Zero Substitution (B8ZS) coding should be used for the TCM. This means a special code is substituted whenever eight consecutive zeros are sent through the link. This code is then interpreted at the remote end of the connection. This technique guarantees ones density independent of the data stream.
hdb3	Indicates that High Density Bipolar (HDB3) coding should be used for the TCM. HDB3 is a bipolar coding method that does not allow more than 3 consecutive zeroes.

6.3.2.2 BITS Framing Configuration

This command allows you to configure the framing format of the DS1 BITS interface or E1 BITS interface for this TCM. You can get to this level by entering **framing** at the **bits** level. Enter the following parameters:

```
myswitch::configuration cec timing bits> framing (d4 | esf) for DS1 interface
              or: framing (fas | fascrc4 | mfas | mfascrc4) for E1 interface
```

The parameters for configuring are defined as follows:

Parameter	Description
d4	Indicates that D4 (also known as Superframe (SF)) framing should be used for the BITS interface on the TCM. SF consists of 12 frames of 192 bits each, with the 193rd bit providing error checking and other functions.
esf	Indicates that Extended Superframe (ESF) framing should be used for the BITS interface on the TCM. ESF provides frame synchronization, cyclic redundancy, and data link bits.
fas	Indicates E1 framing that makes use of the Frame Alignment Signal.
fascrc4	Indicates E1 framing that makes use of the Frame Alignment Signal and CRC-4 checksums.
mfas	Indicates E1 framing that makes use of the Multi-Frame Alignment Signal (sometimes referred to as TS16).
mfascrc4	Indicates E1 framing that makes use of the Multi-Frame Alignment Signal and CRC-4 checksums.

6.3.2.3 BITS Level Configuration

This command lets you configure the output level (dB) of the DS1 BITS interface for this TCM. This command is not valid when using an E1 source.

```
myswitch::configuration cec timing bits> level (0.6 | 1.2 | 1.8 | 2.4 | 3.0)
```

The parameters for configuring are defined as follows:

Parameter	Description
0.6	Indicates that the output level for the BITS interface is 0.6 dB.
1.2	Indicates that the output level for the BITS interface is 1.2 dB.
1.8	Indicates that the output level for the BITS interface is 1.8 dB.
2.4	Indicates that the output level for the BITS interface is 2.4 dB.
3.0	Indicates that the output level for the BITS interface is 3.0 dB.

6.3.3   Failover Timing Configuration

This command lets you configure the timing failover delay. You can display the available sub-command by typing **failover ?** at the **timing** level.

```
myswitch::configuration cec timing> failover ?
  delay
```

6.3.3.1   Configuring the Failover Timing Delay

This command lets you set the delay, in whole seconds, between the time that the primary reference fails and the time that the TCM is told to switch over to the secondary reference. Enter the following parameters:

```
myswitch::configuration cec timing failover> delay <seconds>
```

The parameter for setting is defined as follows:

Parameter	Description
seconds	Indicates the amount of time after the failure of the primary reference before the TCM is instructed to switch to the secondary reference. The default is 0 seconds.

> **NOTE**
>
> The failover delay should be left at 0 under normal circumstances. This parameter is only meaningful when the timing mode is automatic.

6.3.4   Timing Mode Configuration

This command lets you configure the timing reference to be used on the TCM. You can get to this level by entering **mode** at the **timing** level. Enter the following parameters:

```
myswitch::configuration timing> mode (freerun | primary | secondary | automatic | bits)
```

The parameters for configuring are defined as follows:

Parameter	Description
freerun	Indicates that the TCM must use the local oscillator as its timing source.
primary	Indicates that the TCM must use the primary timing source.
secondary	Indicates that the TCM must use the secondary timing source.
automatic	This is the default mode. See the *CEC-Plus Installation and User's Manual* for a description of how this mode works.
bits	Indicates that the TCM will automatically use the most appropriate BITS input. See the *CEC-Plus Installation and User's Manual* for a description of how this mode works.

NOTE

It is recommended that the freerun, primary, and secondary modes only be used during diagnostics and maintenance because the TCM will not failover to another source when it is in one of these modes.

automatic mode is only valid when the primary and secondary sources are recovered from network modules.

6.3.5   Primary Timing Configuration

This command enables you to select the primary clock source from any of the switch fabrics' exported clocks. Enter the following parameters:

```
myswitch::configuration cec timing> primary fabric (1 | 2 | 3 | 4) (primary | secondary)
```

The parameter for configuring is defined as follows:

Parameter	Description
fabric (1 | 2 | 3 | 4) (primary | secondary)	Indicates the primary timing source is to be taken from one of the exported clocks (either primary or secondary) from one of the switch fabrics (1 - first (leftmost slot), 2 - second slot, 3 - third slot, 4 fourth slot) switch fabric.

Before you configure the primary timing source, use the **conf timing references** command to display all of the possible clocking references, their specific sources, and whether or not they are actually available for you to use. The freerun oscillator is not listed there because it is always available.

6.3.6   Displaying Timing References

This commands lets you display all of the possible timing references from the switch. Enter the following parameters:

```
myswitch::configuration cec timing> references
Reference             Source         Status        Admin   RefQual
fabric 1 primary      1A1            available     up      ok
fabric 1 secondary    1 (No Timing)  unavailable   up      ok
fabric 2 primary      2 (Down)       unavailable   up      ok
fabric 2 secondary    2 (Down)       unavailable   up      ok
fabric 3 primary      3 (No Timing)  unavailable   up      ok
fabric 3 secondary    3 (No Timing)  unavailable   up      ok
fabric 4 primary      4 (Down)       unavailable   up      ok
fabric 4 secondary    4 (Down)       unavailable   up      ok
tcm X                 BITS1          unavailable   up      ok
tcm Y                 BITS2          available     up      ok
```

The fields in this display are defined as follows:

Fields	Description
Reference	Shows the name of the timing reference.
Source	Shows from which network module port or BITS clock the timing is being derived for this fabric or TCM. If the port is listed in regular BNP notation (e.g., 3B2), this indicates the exported timing source from this port. If the port is listed as 3 (No Timing), then this network module does not support distributed timing. If the port is listed as 3 OSC, this indicates the timing source is the crystal oscillator on that network module. If the port is listed as 3 (Down), then this network module has been removed.
Status	Shows if the timing source is available or not.
Admin	Shows the administrative status of the timing reference. **up** means it is enabled and **down** means it is disabled.
RefQual	Shows the reference quality (RefQual) status of the timing reference. **ok** means the source is valid. **fail** means the source has gone out of specification. It is recommended that you investigate and correct the cause prior to resetting. To make the source valid, you must manually reset the RefQual status of that source using the **conf cec timing refqual** command.

If no TCM is installed, then the following is displayed:

```
myswitch::configuration cec timing> references
No ESI reference information available
```

6.3.7   Clearing a Reference Quality Failure

This command lets you to clear a reference quality (RefQual) failure for a timing source. Once a RefQual failure has been declared for a timing source, it is considered invalid and cannot be used as a timing source until the failure is cleared. You can use this command to clear the failure and make the source a valid one again.

Before clearing the failure, it is recommended that you investigate the cause of the failure and correct it.

Enter the following:

```
myswitch::configuration cec timing> refqual fabric (1 | 2 | 3 | 4)
(primary | secondary) clear

or

refqual (bits1 | bits2) clear
```

These parameters are defined as follows:

Parameter	Description
fabric (1 | 2 | 3 | 4)	Indicates the switch fabric (1 - first (leftmost slot), 2 - second slot, 3 - third slot, 4 fourth slot) of the timing source for which you are clearing the RefQual failure.
(primary | secondary)	Indicates for which timing source you are clearing the RefQual failure.
(bits1 | bits2)	Indicates for which BITS timing source you are clearing the RefQual failure.

6.3.8   Revertive Timing Configuration

These commands allow you to enable or to disable revertive switching, or to configure the length of the delay between the time that the failed primary clock is restored and the time that the TCM is told to switch back to the primary clock again. You can get to this level by entering **revertive ?** at the **timing** level.

```
myswitch::configuration cec timing> revertive ?
   delay            disable           enable
```

> **NOTE** ➤ Revertive timing is only available when the mode is configured to be **automatic** under **conf timing mode**.

6.3.8.1   Configuring the Revertive Timing Delay

When revertive timing is enabled and the primary clock fails, the TCM is told of the failure and is instructed to switch to the timing source that is configured as the secondary clock. To ensure that the primary source is good when it returns, this command lets you configure the amount of time, in whole seconds, between the time that the primary clock is restored and the time that the TCM is told to switch back to the primary clock again. Enter the following parameters:

```
myswitch::configuration cec timing revertive> delay <seconds>
```

The parameter for configuring is defined as follows:

Parameter	Description
delay	Indicates the amount of time after the restoration of the primary timing reference before the TCM is instructed to return to the primary timing reference. The default is 3 seconds.

6.3.8.2 Disabling Revertive Timing Delay

This command lets you disable the revertive switching for timing sources on this TCM. When you disable revertive switching, you only turn it off. If you enable it again, the TCM uses the value that you last configured for the delay using **conf timing revertive delay**. To disable the revertive timing delay, enter the following parameters:

```
myswitch::configuration cec timing revertive> disable
```

The parameter for disabling is defined as follows:

Parameter	Description
disable	Indicates that the revertive timing delay is going to be turned off.

NOTE ▶ This command takes effect as soon as you enter it.

6.3.8.3 Enabling the Revertive Timing Delay

This command lets you enable the revertive switching delay for timing sources on this TCM. When you enable the delay, the TCM uses the value that you last configured for the delay using **conf timing revertive delay**. To enable the revertive timing delay, enter the following parameters:

```
myswitch::configuration cec timing revertive> enable
```

The parameter for enabling is defined as follows:

Parameter	Description
enable	Indicates that the revertive timing delay is going to be turned on.

NOTE ▶ This command takes effect as soon as you enter it.

6.3.9   Secondary Timing Configuration

This command enables you to select the secondary clock source from any of the switch fabrics' exported clocks. You can get to this level by entering **secondary** at the **timing** level. Enter the following parameters:

```
myswitch::configuration cec timing> secondary fabric (1 | 2 | 3 | 4) (primary|secondary)
```

The parameter for configuring is defined as follows:

Parameter	Description
fabric (1 | 2 | 3 | 4) (primary | secondary)	Indicates the secondary timing source is to be taken from one of the exported clocks (either primary or secondary) from one of the switch fabrics (1 - first (leftmost slot), 2 - second slot, 3 - third slot, 4 fourth slot) switch fabric.

Before you configure the secondary timing source, use the **conf timing references** command to display all of the possible clocking references, their specific sources, and whether or not they are actually available for you to use.

6.3.10 Displaying Timing

This command lets you display the timing information that has been configured. Enter the following parameters:

```
myswitch::configuration cec timing> show
ESI module on 'linus', Card type DS1 Stratum 4
DS1 BITS interface framing esf, line coding b8zs
output level 0.6 dB

PLL Status                        freerun
Current Timing Reference          secondary
Requested Timing Reference        secondary

Primary Reference (2D1)           unavailable
Secondary Reference (2A1)         unavailable

BITS1 Reference                   unavailable
BITS2 Reference                   unavailable

Revertive Switching               enabled
Revertive Switching Delay         10
Failover Switching Delay          3
```

The fields in this display are defined as follows:

Field	Description
ESI module	Shows the name of the TCM.
Card type	Shows the ESI card type for this TCM.
DS1/E1 BITS interface framing	Shows the framing format of the BITS interface for this TCM. For DS1, can be: `d4` or `esf`, and for E1, can be: `fas`, `fascrc4`, `mfas`, or `mfascrc4`.
line coding	Shows the coding format of the BITS interface for this TCM. For DS1, can be: `ami` or `b8zs`, and for E1, can be: `hdb3`.
output level	Shows the output level, in dB, of the DS1 BITS interface for this TCM. Can be `0.6`, `1.2`, `1.8`, `2.4`, or `3.0`.
PLL Status	Shows the phase-locked loop (PLL) status for this TCM. Can be `freerun` meaning the local oscillator on the TCM is being used, `locked` meaning the current clock is good, `holdover` meaning the TCM has detected a clock source error and is using the last valid clock source, `acquire` meaning the TCM is trying to lock on to the current clock (this may take up to five minutes), or `refqual` meaning the reference quality of the new clock is out of specification.
Current Timing Reference	Shows the actual timing reference that is currently being used. Can be `BITS1`, `BITS2`, `primary`, `secondary`, or `freerun`.
Requested Timing Reference	Shows the timing reference that was configured. If this source fails, it will not match what is currently being used. Can be `bits`, `automatic`, `primary`, `secondary`, or `freerun`.
Primary Reference	Shows the timing source configured as the primary source. Can be one of the primary or secondary clocks exported from one of the switch fabrics.
Primary Status	Shows whether the timing source configured as the primary source is currently `available` or `unavailable`.
Secondary Reference	Shows the timing source configured as the secondary source. Can be one of the primary or secondary clocks exported from one of the switch fabrics.
Secondary Status	Shows whether the timing source configured as the secondary source is currently `available` or `unavailable`.
Revertive Switching	Shows whether revertive switching is `enabled` or `disabled`.
Revertive Switching Delay	Shows the amount of time, in seconds, between the time that the primary clock is restored and the time that the TCM switches back to the primary clock again.
Failover Switching Delay	Shows the amount of time, in seconds, after the failure of the primary clock source before the TCM switches to the secondary clock source.

If no TCM is installed, then the following is displayed:

```
myswitch::configuration cec timing> show
No ESI board found
```

AMI Configuration Commands Reference Manual • Part 1

CHAPTER 7   CES Configuration Commands

The `ces` commands let the user create and delete CES connections, as well as display the status of existing connections. You can display the list of available subcommands by typing `?` at the `ces` level.

```
myswitch:: configuration ces> ?
  new              delete          show
```

> **NOTE** These commands are only displayed on the platforms that can support Circuit Emulation (CEM) network modules.

7.1   Creating a New CES Connection

To create a new CES connection, you must configure several parameters. Enter the following to create a new CES connection:

```
myswitch::configuration ces> new <port> <timeslots>
```

The CES **new** command can also be used as shown below. When the following parameters are used, by default, an appropriate entry is made in the UPC table and a bidirectional PVC is created with the proper UPC index.

```
or: new <port> <timeslots> -oport <oport> -ots <ots>
or: new <port> <timeslots> -oport <oport> -ovpi <ovpi> -ovci <ovci>
```

Several advanced options can also be used. For more information about configuring idle channel suppression, see Appendix B in the *ATM Switch Network Configuration Manual*. The following are the advanced options that can be used when creating CES connections:

```
advanced options:
[-srts (on|off)] [-fupc <index>] [-bupc <index>]
[-cas (basic|cas)] [-partialfill <partialfill>] [-reassCDVT <cdvt>]
[-bufSize <bufSize>] [-integ <integ>] [-idlesupp (enable|disable)]
[-idlemask <hex value>] [-idlepat (<idle pattern>|<cas pattern>)]
[-idleintp <msec>]
[-oidlesupp (enable|disable)] [-oidlemask <hex value>]
[-oidlepat (<idle pattern>|<cas pattern>)] [-oidleintp <msec>]
```

> **NOTE** ➤
>
> SRTS is only available on unstructured connections, which are created by specifying **all** for the `<timeslots>` parameter.
>
> The **-cas**, **-partialfill**, **-idlesupp**, and **-oidlesupp**, options do not apply to unstructured mode.
>
> Structured mode is selected by indicating the exact timeslots to be used. For example, timeslots 1, 2, and 3 would be entered as 1-3, timeslots 2, 4, and 6 would be entered as 2:4:6, and combinations such as 1-4:9-11:12 are allowed.

The parameters for new are defined as follows:

Parameter	Description
port	The port on which the CES connection is to be created.
timeslots	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are being configured for a particular PVC. `all` indicates unstructured service. A time slot is a data path of size DS0 in a TDM trunk. The time slot assignments may be either contiguous or non-contiguous DS0s.
oport	The output port of the CES connection, which can be a CES port or an ATM port.
ovpi	The output Virtual Path Identifier (VPI) of the CES connection when the output port is not a CES port.
ovci	The output Virtual Channel Identifier (VCI) of the CES connection when the output port is not a CES port.
ots	The output timeslots of the CES connection when the output port is a CES port.
srts	Indicates whether Synchronous Residual Time Stamp (SRTS) clock recovery is to be enabled on this connection. `on` indicates that SRTS is enabled, `off` indicates that SRTS is disabled. The default is `off`.
-fupc <index>	The UPC contract type to be used in the ingress direction of the connection. (See Part 2 of the *AMI Configuration Commands Reference Manual* for more information about UPC contracts.)
-bupc <index>	The UPC contract type to be used in the egress direction of the connection. (See Part 2 of the *AMI Configuration Commands Reference Manual* for more information about UPC contracts.)
cas	Indicates whether Channel Associated Signalling (CAS) is to be used on the connection. `basic` indicates that CAS will not be used, `cas` indicates that CAS will be used. The default is `basic`.
partialfill	Indicates how many of the available 47 payload bytes in each cell are used before they are deemed "full" and ready for transmission across the ATM network (i.e., how much of the ATM cell contains data and how much is filler). The range for this parameter is 12 to 47. The default value is `47`, for 47 bytes of data. On structured basic connections, a minimum partialfill size of 24 is required when setting the idle channel suppression (`-idlesupp`) option to enabled. partialfill is used to minimize network transmission latency and is useful especially with time-sensitive, robbed-bit signalling sources.
-reassCDVT <cdvt>	The Cell Delay Variation Tolerance for cells being received by the segmentation and reassembly (SAR) engine. The range for this parameter is 100 to 24000 (in μs), and the default is `2000` (i.e., 2 μs).
bufSize	The amount of reassembly buffer space allocated for the connection. The default is `256` bytes per timeslot.
integ	The amount of time allocated to re-establish the connection before, while, or after the call is established, or in the case of interruption. The default is `2500` μs.

CES Configuration Commands

Parameter	Description
-idlesupp (enable | disable)	**enable** means idle channel suppression is going to be used on the incoming side of this connection. **disable** means idle channel suppression is not going to be used on the incoming side of this connection. The default is **disable**. When enabling this option on structured basic connections, a minimum partialfill size of 24 is required.
-idlemask <hex value>	The mask pattern for idle detection on the incoming side of this connection. This method is only used for structured basic service. The range of values is 01 to FF. The default is **FF**.
-idlepat (<idle pattern> | <cas pattern>)	The pattern for idle detection on the incoming side of this connection. For detection based on both idle and mask patterns, it contains the idle octet pattern. For detection based on signalling, it contains one of the following CAS patterns: em00, em01, fxolsuser01, fxolsuser11, fxolsnet00, fxolsnet01, fxslsuser00, fxslsuser01, fxslsnet01, fxslsnet11, fxsgsuser01, fxsgsnet10, fxsgsnet11, fxogsuser10, fxogsuser11, fxogsnet01, r210. A maximum of one idle pattern can be used for structured basic connections. Idle patterns are filled from the least significant byte. The default is **FF** for structured basic connections. The default CAS idle AB bit pattern is **em00**.
-idleintp <msec>	The integration period for idle detection on the incoming side of this connection, in msec. Idle patterns are observed for this period before declaring that an active connection has gone idle. For both CAS and basic connections, the range of values is 500 μs to 2 seconds and the default is **1** second.
-oidlesupp (enable | disable)	**enable** means idle channel suppression is going to be used on the outgoing side of this connection. **disable** means idle channel suppression is not going to be used on the outgoing side of this connection. The default is **disable**. When enabling this option on structured basic connections, a minimum partialfill size of 24 is required.
-oidlemask <hex value>	The mask pattern for idle detection on the outgoing side of this connection. This method is only used for structured basic service. The range of values is 01 to FF. The default is **FF**.
-oidlepat (<idle pattern> | <cas pattern>)	The pattern for idle detection on the outgoing side of this connection. For detection based on both idle and mask patterns, it contains the idle octet pattern. For detection based on signalling, it contains one of the following CAS patterns: em00, em01, fxolsuser01, fxolsuser11, fxolsnet00, fxolsnet01, fxslsuser00, fxslsuser01, fxslsnet01, fxslsnet11, fxsgsuser01, fxsgsnet10, fxsgsnet11, fxogsuser10, fxogsuser11, fxogsnet01, r210. A maximum of one idle pattern can be used for structured basic connections. Idle patterns are filled from the least significant byte. The default is **FF** for structured basic connections. The default CAS idle AB bit pattern is **em00**.
-oidleintp <msec>	The integration period for idle detection on the outgoing side of this connection, in msec. Idle patterns are observed for this period before declaring that an active connection has gone idle. For both CAS and basic connections, the range of values is 500 μs to 2 seconds and the default is **1** second.

7.2   Deleting a CES Connection

To delete a CES connection, enter the following parameters:

```
myswitch::configuration ces> delete <service>
```

The parameter for delete is defined as follows:

Parameter	Description
service	The CES service ID of the connection to be deleted.

7.3   Displaying CES Connections

To display the current CES connections, enter the following:

```
myswitch::configuration ces> show
Searching For CES ports.....
  CES         Input                                    Output
  Service State Port  Timeslots     VPI VCI  Type Port TimeSlots     VPI VCI
  24      down  1A1   1             0   129  -    -    -             -   -
  31      down  1A1   2-3           0   130  spvc 1D3  -             0   32
  32      down  1A1   4-5           0   131  pvc  1D4  -             0   150
  33      down  1A1   6-7           0   132  spvc 1D3  -             0   35
```

The fields in this display are defined as follows：

Field	Description
CES Service	The identification number (assigned by the switch) of this CES connection.
State	Indicates whether the CES connection is enabled (up) or disabled (down).
Input Port	The incoming port on which the CES connection exists.
Timeslots	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are configured for the input port. **all** indicates unstructured service.
Input VPI	The incoming VPI value of the connection.
Input VCI	The incoming VCI value of the connection.
Type	The type of ATM connection (i.e., PVC or SPVC) that is associated with the CES connection.
Output Port	The outgoing port on which the CES connection exists.
Timeslots	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are configured for the output port. **all** indicates unstructured service.
Output VPI	The outgoing VPI value of the connection.
Output VCI	The outgoing VCI value of the connection.

To display the advanced settings of the current CES connections, enter the following:

```
myswitch::configuration ces> show advanced
Searching For CES ports.....
                    Service    Clock          Partial    Max        Integ.
Service MapVPI MapVCI Type     Mode  Cas  Fill   BufSize   CDVT  Period
2024    0      129    structured synch basic 0           256       900   2500
```

The fields in this display are defined as follows：

Field	Description
CES Service	The identification number (assigned by the switch) of this CES connection.
MapVPI	The incoming VPI value of the connection.
MapVCI	The incoming VCI value of the connection.
Service Type	Shows if this connection uses structured or unstructured service.
Clock Mode	`Synch` means that the connection is in synchronous mode (either structured or unstructured). `SRTS` means that the connection is in asynchronous (unstructured) mode. (Synchronous Residual Time Stamp (SRTS) clock recovery is enabled on this connection.)
Cas	`basic` indicates that Channel Associated Signalling (CAS) will not be used, `cas` indicates that CAS will be used.
Partial Fill	Indicates how many of the available 47 payload bytes in each cell are used before they are deemed "full" and ready for transmission across the ATM network (i.e., how much of the ATM cell contains data and how much is filler). partialfill is used to minimize network transmission latency and is useful especially with time-sensitive, robbed-bit signalling sources.
Max BufSize	The amount of reassembly buffer space allocated for the connection. The default is 256 bytes per timeslot.
CDVT	The Cell Delay Variation Tolerance for cells being received by the segmentation and reassembly (SAR) engine. The range for this parameter is 100 to 24000 (in μs), and the default is 2000 (i.e., 2 μs).
Integ. Period	The amount of time allocated to re-establish the connection before, while, or after the call is established, or in the case of interruption. The default is 2500μs.

CES Configuration Commands

To display the idle channel suppression information for the current CES connections, enter the following:

```
myswitch::configuration ces> show [<id>] [(advanced|idlesupp)]
myswitch::configuration ces> show idlesupp
CES     Idlesupp  Idle     Idle  NoOf  Idle      IdleInt.
Service State     Type     Mask  Patt  Patterns  Period(µs)
20024   disabled  -        -     -     -         -
20031   enabled   pattern  FF    1     7F        500
20032   enabled   cas      01b   1     10b       100
20033   disabled  -        -     -     -         -
20034   disabled  -        -     -     -         -
20035   enabled   pattern  FF    1     7F        500
20036   enabled   cas      11    1     00        100
20037   disabled  -        -     -     -         -
```

The fields in this display are defined as follows：

Field	Description
CES Service	The identification number (assigned by the switch) of this CES connection.
Idlesupp State	**enabled** means idle channel suppression is going to be used on this connection. **disabled** means idle channel suppression is not going to be used on this connection.
Idle Type	The idle detection mechanism used on this connection. **pattern** means idle detection is based on idle patterns. **cas** means idle detection is based on CAS signalling. This field is read-only.
Idle Mask	The mask pattern for idle detection on this connection. This method of idle detection can only be used on basic and unstructured connections.
NoOf Patt	The number of idle patterns configured for detection. This field is read-only.
Idle Patterns	The patterns for idle detection on this connection. For detection based on both idle and mask patterns, it contains the idle octet patterns. For detection based on signalling, it contains one of the following patterns: em00, em01, fxolsuser01, fxolsuser11, fxolsnet00, fxolsnet01, fxslsuser00, fxslsuser01, fxslsnet01, fxslsnet11, fxsgsuser01, fxsgsnet10, fxsgsnet11, fxogsuser10, fxogsuser11, fxogsnet01, r210. A maximum of one idle pattern can be used for structured basic connections. Idle patterns are filled from the least significant byte.
IdleInt. Period (µs)	The integration period for idle detection on this connection. Idle patterns are observed for this period before declaring that an active connection has gone idle.

If no CES network modules are installed, then the following is displayed:

```
myswitch::configuration ces> show
Searching For CES ports.....
No services found
```

CHAPTER 8   Frame Relay Commands

These commands let you configure Frame Relay services and PVCs on a *FramePlus* network module. These commands are only displayed on the platforms that can support *FramePlus* network modules. Type **?** at the **fratm** level to display the available subcommands.

```
myswitch::configuration fratm> ?
  delete           new              pvc>             admin
  egress_re        stats            trap             show
```

8.1   Deleting a Frame Relay Service

This command lets you delete a Frame Relay service. Enter the following:

```
myswitch::configuration fratm> delete <serviceId>
```

For example, to delete a service, enter something similar to the following:

```
myswitch::configuration fratm> del 2A1:00
```

The parameter for delete is defined as follows:

Parameter	Description
serviceId	The ID of the Frame Relay service to be deleted. Look under the SvcId field under **conf fratm show** to find this number.

> **NOTE**
>
> Before you can delete a service, you must delete all connections that are using that service. To see how many connections are using a particular service, look at the Vccs field under **stats fratm service**.

8.2   Creating a Frame Relay Service

This command lets you create a Frame Relay service. Enter the following:

NOTE ► You must either use the default Frame Relay profiles or create new ones using the **conf profile** commands and then create the Frame Relay services using this command. Then, you can create a Frame Relay PVC using the **conf fratm pvc** commands. See Appendix D in the *ATM Switch Network Configuration Manual* for your switch for more information about this process.

```
myswitch::configuration fratm> new <port> <timeslots> [-lmi <index>] [-service <index>]
        [-egress_re (enable|disable)] [-admin (up|down)]
        [-name <name>] [-id <bnp:cc>]
```

NOTE ► Timeslots should be entered in one of the following formats: Timeslots 1, 2, and 3 should be entered as 1-3. Timeslots 2, 4, and 6 should be entered as 2:4:6. Combinations such as 1-4:9-11:13 are allowed. Unchannelized connections should be entered as **all** (this means 1-24 for DS1 and 1-31 for E1).

If you wish to use multiple timeslots, you must ensure that the access rate of the referenced service profile is large enough to support them. The default access rate is 64 Kbps. To verify the access rate, look in the accRate field under **conf profile service**. If the rate is 64, you can only use a single timeslot. If it says 128 Kbps, you can use two timeslots. If it says 1536 Kbps, you can use 24 timeslots, etc.

For example, to create a Frame Relay service, enter something similar to the following:

```
myswitch::configuration fratm> new 1a1 1 -lmi 3 -service 2 -name service_a
The newly created service id is 1A1:00
```

As shown above, the switch responds with a service ID, which can be referenced in other AMI commands that prompt for a `<serviceid>` such as `conf fratm pvc new`. If you do not specify a service ID at the time of creation, a service ID is assigned automatically.

The parameters for new are defined as follows:

Parameter	Description
port	The port on which the Frame Relay service is to be created.
timeslots	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are being configured for this service. `all` indicates unchannelized connections. The time slot assignments must match the access rate specification in the service profile. (Look in the `accRate` field under `conf profile service`.) The access rate for one timeslot is 64 Kbps, and the access rate for two timeslots is 128 Kbps, etc.
lmi <index>	The lmi profile that is to be used with this service. Look under the `Index` field under `conf profile lmi show` to find this number.
service <index>	The service profile that is to be used with this service. Look under the `Index` field under `conf profile service show` to find this number.
egress_re (enable | disable)	Enables or disables egress rate enforcement for this service. It ensures that frames transmitted on the DS1 or E1 interface do not violate any traffic contract defined for a particular connection being supported by this service. Because the traffic is managed at a connection level, the rate enforcement algorithm uses the values of the outBc, outBe, and outCir arguments within the connection's frrate profile. (See those respective fields under `conf profile frrate show`.) `enabled` means rate enforcement is active. `disabled` means rate enforcement is not active. The default value is `disabled`.
admin (up | down)	Sets the administrative status of this service. `enabled` means it is up. `disabled` means it is down. The default value is `enabled`.
name <name>	The optional, user-assigned name for this service.
id <bnp:cc>	The optional, user-assigned service identifier (ID) for this service. This service ID can be referenced in other AMI commands that prompt for a `<serviceid>` such as `conf fratm pvc new`. The bnp portion of the service ID must be identical to the port number that was entered for `<port>`. The cc portion of the service ID must be between 0 and 23 for a DS1 interface and 0 and 30 for an E1 interface. If you do not enter a service ID, the switch assigns one for you and displays it as soon as you enter the complete command.

Frame Relay Commands

8.3   Configuring Frame Relay PVCs

These commands let you configure Frame Relay PVCs. Type **?** at the **pvc** level to display the available subcommands.

```
myswitch::configuration fratm pvc> ?
  delete              new              admin            ingress_re
  show
```

8.3.1   Deleting a Frame Relay PVC

This command lets you delete a Frame Relay PVC. Enter the following:

```
myswitch::configuration fratm pvc> delete <serviceId> <dlci>
```

For example, to delete a Frame Relay PVC, enter something similar to the following:

```
myswitch::configuration fratm pvc> del 4a1:01 101
```

8.3.2   Creating a Frame Relay PVC

This command lets you create a Frame Relay PVC.

> **NOTE**
>
> You must either use the default Frame Relay profiles or create new ones using the **conf profile** commands and then create the Frame Relay services using the **conf fratm new** command. Then, you can create a Frame Relay PVC using this command. See Appendix D in the *ATM Switch Network Configuration Manual* for your switch for more information about this process.

> **NOTE**
>
> Currently, the *FramePlus* network module does not permit the establishment of Frame-to-Frame connections across a single fabric. This restriction prevents connections between FramePlus ports located on the same network module, or *FramePlus* ports located across separate modules, but attached to the same switch fabric.

Enter the following:

```
myswitch::configuration fratm pvc> new <serviceid> <dlci> [-oport <oport>]
[-ovpi <ovpi>]
[-ovci <ovci>] [-faupc <index>] [-afupc <index>] [-epdppd <index>]
[-admin (up|down)] [-frrate <index>] [-frf8 <index>]
[-name <name>]

note: if oport/ovpi/ovci aren't specified, a dangling FRATM PVC will be created.
```

The parameters for delete and new are defined as follows:

Parameter	Description
serviceId	The ID of the service to be used by this PVC. Look in the `SvcId` field under **conf fratm show** to find this number.
dlci	The Data Link Connection Identifier used to identify this Frame Relay PVC.
oport	The output port for this Frame Relay PVC, which is an ATM port.
ovpi	The output Virtual Path Identifier (VPI) of this Frame Relay PVC.
ovci	The output Virtual Channel Identifier (VCI) of this Frame Relay PVC.
faupc <index>	The UPC traffic contract that is to be used for this PVC. Look under the `Index` field under **conf upc show** to find this number. This contract is applied to the connection in the Frame Relay to ATM direction.
afupc <index>	The UPC traffic contract that is to be used for this PVC. Look under the `Index` field under **conf upc show** to find this number. This contract is applied to the connection in the ATM to Frame Relay direction.
epdppd <index>	The EPD/PPD profile that is to be used for this PVC. Look under the `Index` field under **conf profile epdppd show** to find this number.
admin (up | down)	Changes the administrative status of this PVC. **up** means it is enabled. **down** means it is disabled. The default value is **up**.
frrate <index>	The Frame Relay rate profile that is to be used for this PVC. Look under the `Index` field under **conf profile frrate show** to find this number.
frf8 <index>	The FRF.8 profile that is to be used for this PVC. Look under the `Index` field under **conf profile frf8 show** to find this number.
name <name>	The optional, user-assigned name for this PVC.

**Frame Relay
Commands**

8.3.3   Changing the Status of a Frame Relay PVC

This command lets you configure the administrative status of a particular Frame Relay PVC on a *FramePlus* network module to be up or down. Enter the following:

```
myswitch::configuration fratm pvc> admin <serviceid> <dlci> (up|down)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The ID of the Frame Relay PVC whose administrative status needs to be changed. The ID number can be found in the SvcId field under **conf fratm pvc show**.
dlci	The Data Link Connection Identifier that indicates the particular Frame Relay PVC whose administrative status needs to be changed. The DLCI can be found in the dlci field under **conf fratm pvc show**.
up | down	**up** sets the administrative status of the designated PVC to up (enabled). **down** sets the administrative status of the designated PVC to down (disabled).

8.3.4   Setting Ingress Rate Enforcement on a Frame Relay PVC

This command lets you enable or disable ingress rate enforcement on an existing Frame Relay PVC on a *FramePlus* network module. Enter the following:

```
myswitch::configuration fratm pvc> ingress_re <serviceid> <dlci> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The ID of the Frame Relay service. The ID number can be found in the SvcId field under **conf fratm pvc show**.
dlci	The Data Link Connection Identifier that indicates the particular Frame Relay PVC that is to be configured. The DLCI can be found in the dlci field under **conf fratm pvc show**.
disable | enable	**enable** means ingress rate enforcement is used on this Frame Relay PVC. **disable** means standard ingress rate enforcement is not used on this Frame Relay PVC. The default is **disable**.

8.3.5   Displaying Frame Relay PVC Information

This command lets you display information about Frame Relay PVCs. Enter the following:

```
myswitch::configuration fratm pvc> show
             Input         Output
 SvcId  dlci Port  VPI VCI Port VPI VCI Admin   Epd Frr Frf8  IngRE    Name
 4A1:00  40  4A1   0   32  4C1   0   40   up     0   0   0   disable  pvc_a
 4A1:00  40  4C1   0   40  4A1   0   32   up     0   0   0   disable  pvc_a
 4A1:01  41  4A1   16  32  4C1   0   41   up     0   0   0   enable   pvc_b
 4A1:01  41  4C1   0   41  4A1   16  32   up     0   0   0   enable   pvc_b
 4A1:02  42  4A1   32  32  4C1   0   42   up     0   0   0   disable  pvc_c
 4A1:02  42  4C1   0   42  4A1   32  32   up     0   0   0   disable  pvc_c
```

The fields in this display are defined as follows：

Field	Description
SvcId	The ID of the service that is used by this PVC.
dlci	The Data Link Connection Identifier used to identify this Frame Relay PVC.
Input Port	The input port of the Frame Relay PVC.
Input VPI	The input Virtual Path Identifier (VPI) of the Frame Relay PVC.
Input VCI	The input Virtual Channel Identifier (VCI) of the Frame Relay PVC.
Output Port	The output port of the Frame Relay PVC.
Output VPI	The output Virtual Path Identifier (VPI) of the Frame Relay PVC.
Output VCI	The output Virtual Channel Identifier (VCI) of the Frame Relay PVC.
Admin[1]	The administrative state of this PVC. `up` means it is enabled. `down` means it is disabled.
Epd	The index of the EPD/PPD profile that is used by this PVC.
Frr	The index of the Frame Relay rate profile that is used by this PVC.
Frf8	The index of the FRF.8 profile that is used by this PVC.
IngRE	`dis` (disabled) means standard ingress rate enforcement is not used by this Frame Relay PVC. `enb` (enabled) means ingress rate enforcement is used by this Frame Relay PVC.
Name	The user-assigned name for this PVC.

[1.] The admin state or status under the `conf port iwf ds1/e1`, `conf fratm`, `conf fratm pvc`, `conf funi` and `conf funi pvc` levels reflects the configured state of a port, service, or connection. The admin status is unique to a level. For example, if a connection under `conf fratm pvc` is administered down, the corresponding admin state of the service under `conf fratm` is up, but the connection will not pass traffic. This is because the admin status of each of these entities has a direct effect on the operational status of all three entities. The operational status of these entities is currently not displayed in AMI.

You can also display information about a particular service as follows:

```
myswitch::configuration fratm pvc> show ?
usage: show [<serviceId>] [<dlci>] [advanced]
myswitch::configuration fratm pvc> show 4a1:01
            Input         Output
 SvcId  dlci Port  VPI VCI  Port VPI VCI Admin  Epd Frr Frf8 IngRE Name
 4A1:01   41   4A1  16  32   4C1   0  41    up    0   0    0   dis  pvc_b
 4A1:01   41   4C1   0  41   4A1  16  32    up    0   0    0   dis  pvc_b
```

The fields in this display are defined in the same manner as those in the previous example.

You can display information about a particular service and DLCI as follows:

```
myswitch::configuration fratm pvc> show 4a1:01 41
            Input         Output
 SvcId  dlci Port  VPI VCI  Port VPI VCI Admin  Epd Frr Frf8 IngRE Name
 4A1:01   41   4A1  16  32   4C1   0  41    up    0   0    0   dis  pvc_b
 4A1:01   41   4C1   0  41   4A1  16  32    up    0   0    0   dis  pvc_b
```

The fields in this display are defined in the same manner as those in the previous example.

You can also display advanced information as follows:

```
myswitch::configuration fratm pvc> show advanced
 SvcId  dlci afupc faupc
 4A1:00   40    0     0
 4A1:01   41    0     0
 4A1:02   42    0     0
```

The fields in this display are defined as follows：

Field	Description
SvcId	The ID of the service that is used by this PVC.
dlci	The Data Link Connection Identifier used to identify this Frame Relay PVC.
afupc	The UPC traffic contract that is used on this PVC in the ATM to Frame Relay direction.
faupc	The UPC traffic contract that is used on this PVC in the Frame Relay to ATM direction.

If no *FramePlus* network modules are installed, or if there are no Frame Relay PVCs config-
ured, then the following is displayed:

```
myswitch::configuration fratm pvc> show
No FR/ATM VCC information available
```

8.4 Changing the Status of a Frame Relay Service

This command lets you configure the administrative status of a particular Frame Relay service on a *FramePlus* network module to be up or down. Enter the following:

```
myswitch::configuration fratm> admin <serviceid> (up|down)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The Frame Relay service whose administrative status needs to be changed. The ID number can be found in the SvcId field under **conf fratm show**.
up | down	**up** changes the administrative status of the designated service to up (enabled). **down** changes the administrative status of the designated service to down (disabled).

8.5 Configuring Egress Rate Enforcement

This command enables you to to enable or disable egress rate enforcement on an existing Frame Relay service. Enter the following:

```
myswitch::configuration fratm> egress_re <serviceid> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The Frame Relay service that is to be configured. The ID number can be found in the SvcId field under **conf fratm show**.
enable | disable	Changes the status of egress rate enforcement for this service. It ensures that frames transmitted on the DS1 or E1 interface do not violate any traffic contract defined for a particular connection being supported within this service. Because the traffic is managed at a connection level, the rate enforcement algorithm is determined by the values of the outBc, outBe, and outCir arguments within the connection's frrate profile. (See those respective fields under **conf profile frrate show**.) **enable** means rate enforcement is active. **disable** means rate enforcement is not active. The default value is **disable**.

8.6   Enabling and Disabling FR Service Statistics

This command allows you to enable or disable the collection of statistics on an existing Frame Relay service.

> **NOTE**
>
> To collect Frame Relay statistics, they must be enabled both here <u>and</u> at the network module level under **conf module fram stats**. By default, they are disabled here, but enabled at the module level.

Enter the following:

```
myswitch::configuration fratm> stats <serviceid> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The Frame Relay service that is to be configured. The ID number can be found in the SvcId field under **conf fratm show**.
enable | disable	**enable** means statistics are collected on the designated service. **disable** means statistics are not collected on the designated service. The default is **disable**.

8.7   Enabling and Disabling FR Service Traps

This command allows you to to enable or disable the SNMP service UP and DOWN traps on an existing Frame Relay service. Enter the following:

```
myswitch::configuration fratm> trap <serviceid> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The Frame Relay service that is to be configured. The ID number can be found in the SvcId field under **conf fratm show**.
enable | disable	**enable** means SNMP UP and DOWN traps are operational on the designated service. **disable** means SNMP UP and DOWN traps are not operational on the designated service. The default is **enable**.

8.8 Displaying Frame Relay Service Information

To display information about the current Frame Relay services, enter the following:

```
myswitch::configuration fratm> show
Searching for FR-ATM services
 SvcId  TimeSlot     Admin EgressRE  Lmi Serv  Traps    Stats     Name
 4A1:00 1            up    disabled  0   0      enabled  disabled  service_a
 4A1:01 2            up    disabled  0   0      enabled  disabled  service_b
 4A1:02 3            up    disabled  0   0      enabled  disabled  service_c
```

The fields in this display are defined as follows：

Field	Description
SvcId	The ID for this service.
TimeSlot	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are configured for this service. **all** indicates unchannelized connections. The time slot assignments must match the access rate specification in the service profile.
Admin[1]	The administrative status of this service. **up** means it is enabled. **down** means it is disabled.
EgressRE	The status of egress rate enforcement on this service. **enabled** means it is on. **disabled** means it is off.
Lmi	The index of the LMI profile that is used by this service.
Serv	The index of the service profile that is used by this service.
Traps	**enabled** means SNMP traps are being generated for this service. **disabled** SNMP traps are not being generated for this service.
Stats	**enabled** means statistics are being collected for this service. **disabled** statistics are not being collected for this service.
Name	The optional, user-assigned name for this service.

[1.] The admin state or status under the **conf port iwf ds1/e1**, **conf fratm**, **conf fratm pvc**, **conf funi** and **conf funi pvc** levels reflects the configured state of a port, service, or connection. The admin status is unique to a level. For example, if a connection under **conf fratm pvc** is administered down, the corresponding admin state of the service under **conf fratm** is up, but the connection will not pass traffic. This is because the admin status of each of these entities has a direct effect on the operational status of all three entities. The operational status of these entities is currently not displayed in AMI.

You can also display information about a particular service as follows:

```
myswitch::configuration fratm> show [<serviceId>]
myswitch::configuration fratm> show 4a1:00
Searching for FR-ATM services
 SvcId  TimeSlot    Admin EgressRE  Lmi Serv Traps    Stats    Name
 4A1:00 1           up    disabled  0   0    enabled  disabled service_a
```

The fields in this display are defined in the same manner as those in the previous example.

If no *FramePlus* network modules are installed or if no services are configured, then the following is displayed:

```
myswitch::configuration fratm> show
Searching for FR-ATM services
....No services found
```

CHAPTER 9 **FUNI Configuration Commands**

These commands let you configure FUNI services. These commands are only displayed on the platforms that can support *FramePlus* network modules. Type **?** at the **funi** level to display the available subcommands.

```
myswitch::configuration funi> ?
  delete            new             pvc>              admin
  stats             trap            show
```

9.1   Deleting a FUNI Service

This command lets you delete a FUNI service. Enter the following parameters:

```
myswitch::configuration funi> delete <serviceId>
```

For example, to delete a FUNI service, enter something similar to the following:

```
myswitch::configuration funi> del 2A1:00
```

The parameter for delete is defined as follows:

Parameter	Description
serviceId	The ID of the FUNI service to be deleted. Look in the **SvcId** field under **conf funi show** to find this number.

> **NOTE**
>
> Before you can delete a service, you must delete all connections that are using that service. To see how many connections are using a particular service, look at the Vccs field under **stats funi service**.

9.2 Creating a FUNI Service

This command lets you create a FUNI service. Enter the following:

NOTE You must either use the default profiles or create FUNI profiles using the **conf profile** commands and then create the FUNI services using this command. Then, you can create a FUNI PVC using the **conf funi pvc** commands. See Appendix D in the *ATM Switch Network Configuration Manual* for your switch for more information about this process.

```
myswitch::configuration funi> new <port> <timeslots> [-funi <index>] [-service <index>]
         [-admin (up|down)] [-name <name>] [-id <bnp:cc>]
```

NOTE Timeslots should be entered in one of the following formats: Timeslots 1, 2, and 3 should be entered as 1-3. Timeslots 2, 4, and 6 should be entered as 2:4:6. Combinations such as 1-4:9-11:13 are allowed. Unchannelized connections should be entered as **all** (this means 1-24 for DS1 and 1-31 for E1).

If you wish to use multiple timeslots, you must ensure that the access rate of the referenced service profile is large enough to support them. The default access rate is 64 Kbps. To verify the access rate, look in the accRate field under **conf profile service**. If the rate is 64, you can only use a single timeslot. If it says 128 Kbps, you can use two timeslots. If it says 1536 Kbps, you can use 24 timeslots, etc.

For example, to create a FUNI service, enter something similar to the following:

```
myswitch::configuration funi> new 4a1 1 -funi 1 -name service_a
The newly created service id is 4A1:00
```

As shown above, the switch responds with a service ID, which can be referenced in other AMI commands that prompt for a `<serviceid>` such as **conf funi pvc new**. If you do not specify a service ID at the time of creation, a service ID is assigned automatically.

The parameters for new are defined as follows:

Parameter	Description
port	The port on which the FUNI service is to be created.
timeslots	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are being configured for this service. **all** indicates unchannelized connections. The time slot assignments must match the access rate specification in the service profile. (Look in the accRate field under **conf profile service**.) The access rate for one timeslot is 64 Kbps, and the access rate for two timeslots is 128 Kbps, etc.
funi <index>	The FUNI profile that is to be used for this service. Look under the Index field under **conf profile funi show** to find this number.
service <index>	The generic service profile that is to be used for this service. Look under the Index field under **conf profile service show** to find this number.
admin (up | down)	Sets the administrative status of this service. **up** means it is enabled. **down** means it is disabled.
name <name>	The optional, user-assigned name for this service.
id <bnp:cc>	The optional, user-assigned service identifier (ID) for this service. This service ID can be referenced in other AMI commands that prompt for a `<serviceid>` such as **conf funi pvc new**. The bnp portion of the service ID must be identical to the port number that was entered for `<port>`. The cc portion of the service ID must be between 0 and 23 for a DS1 interface and 0 and 30 for an E1 interface. If you do not enter a service ID, the switch assigns one for you and displays it as soon as you enter the complete command.

9.3   Configuring FUNI PVCs

This command lets you configure FUNI PVCs. Type **?** at the **pvc** level to display the available subcommands.

```
myswitch::configuration funi pvc> ?
  delete              new              admin              show
```

9.3.1   Deleting a FUNI PVC

This command lets you delete a FUNI PVC. Enter the following:

```
myswitch::configuration funi pvc> delete <serviceId> <fvpi> <fvci>
```

For example, to delete a FUNI PVC, enter something similar to the following:

```
myswitch::configuration funi pvc> del 4d1:01 0 44
```

9.3.2   Creating a FUNI PVC

This command lets you create a FUNI PVC. Enter the following:

NOTE ➤ You must first create any profiles that you want using the **conf profile** commands and then create the FUNI services using the **conf funi new** command. Then, you can create a FUNI PVC using this command. See Appendix D in the *ATM Switch Network Configuration Manual* for your switch for more information about this process.

```
myswitch::configuration funi pvc> new <serviceid> <fvpi> <fvci> [-oport <oport>]
[-ovpi <ovpi>] [-ovci <ovci>] [-faupc <index>] [-afupc <index>] [-epdppd <index>]
[-admin (up|down)] [-name <name>]

note: if oport/ovpi/ovci aren't specified, a dangling FRATM PVC will be created.
```

For example, to create a FUNI PVC, enter something similar to the following:

```
myswitch::configuration funi pvc> new 4A1:00 0 40 -oport 4C1 -fvpi 0 -fvci 40 -name pvc_a
```

The parameters for delete and new are defined as follows:

Parameter	Description
serviceId	The ID of the service to be used by this PVC. Look under the `SvcId` field under **conf funi show** to find this number.
fvpi	The Virtual Path Identifier (VPI) of the FUNI PVC. This number must be within the range specified by the FUNI profile that is associated with the service used. The number of active bits is displayed in the `#bits Vpi` field under **conf profile funi show**.
fvci	The Virtual Channel Identifier (VCI) of the FUNI PVC. This number must be within the range specified by the FUNI profile that is associated with the service used. The range is displayed in the `vci Min` and `vci Max` fields under **conf profile funi show**.
oport	The output port of the FUNI PVC.
ovpi	The output Virtual Path Identifier (VPI) of the FUNI PVC on the ATM side.
ovci	The output Virtual Channel Identifier (VCI) of the FUNI PVC ATM.
faupc <index>	The UPC traffic contract that is to be used for this PVC. Look under the `Index` field under **conf upc show** to find this number. This contract is applied to the PVC in the FUNI to ATM direction.
afupc <index>	The UPC traffic contract that is to be used for this PVC. Look under the `Index` field under **conf upc show** to find this number. This contract is applied to the PVC in the ATM to FUNI direction.
epdppd <index>	The EPD/PPD profile that is to be used for this PVC. Look under the `Index` field under **conf profile epdppd show** to find this number.
admin (up | down)	Changes the administrative status of this PVC. **up** means it is enabled. **down** means it is disabled. The default value is **up**.
name <name>	The optional, user-assigned name for this connection.

9.3.3   Changing the Status of a FUNI PVC

This command enables you to configure the administrative status of a particular FUNI PVC on a *FramePlus* network module to be up or down. Enter the following:

```
myswitch::configuration funi pvc> admin <serviceid> <fvpi> <fvci> (up|down)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The ID of the FUNI service. The ID number can be found in the `SvcId` field under **conf funi pvc show**.
fvpi	The Virtual Path Identifier (VPI) of the FUNI PVC.
fvci	The Virtual Channel Identifier (VCI) of the FUNI PVC.
up | down	**up** sets the administrative status of the designated service to up (enabled). **down** sets the administrative status of the designated service to down (disabled).

9.3.4   Displaying FUNI PVC Information

This command lets you display information about FUNI PVCs. Enter the following:

```
myswitch::configuration funi pvc> show
        FUNI FUNI Input         Output
 SvcId   VPI  VCI  Port VPI VCI  Port VPI VCI  Admin Eppd Name
 4A1:02    0   42   4A1  32  42   4C1   0  42    up    0 pvc_c
 4A1:02    0   42   4C1   0  42   4A1  32  42    up    0 pvc_c
 4A1:00    0   40   4A1   0  40   4C1   0  40    up    0 pvc_a
 4A1:00    0   40   4C1   0  40   4A1   0  40    up    0 pvc_a
 4A1:01    0   41   4A1  16  41   4C1   0  41    up    0 pvc_b
 4A1:01    0   41   4C1   0  41   4A1  16  41    up    0 pvc_b
```

The fields in this display are defined as follows:

Field	Description
SvcId	The ID of the service used by this PVC.
FUNI VPI	The Virtual Path Identifier (VPI) of the PVC on the FUNI side that you entered for `<fvpi>`. This value may or may not be used by the switch software. The value that is actually used is displayed in the `Input VPI` field.
FUNI VCI	The Virtual Channel Identifier (VCI) of the PVC on the FUNI side that you entered for `<fvci>`. This value may or may not be used by the switch software. The value that is actually used is displayed in the `Input VCI` field.
Input Port	The input port of the PVC on the FUNI side.
Input VPI	The input Virtual Path Identifier (VPI) of the FUNI PVC that is selected and used by the switch software. This value may or may not be what you entered for `<fvpi>`.
Input VCI	The input Virtual Channel Identifier (VCI) of the FUNI PVC that is selected and used by the switch software. This value may or may not be what you entered for `<fvci>`.
Output Port	The output port of the PVC on the ATM side.
Output VPI	The output Virtual Path Identifier (VPI) of the PVC on the ATM side.
Output VCI	The output Virtual Channel Identifier (VCI) of the PVC on the ATM side.
Admin[1]	The administrative status of this PVC. **enabled** means it is up. **disabled** means it is down.
Eppd	The EPD/PPD profile that is used by this PVC.
Name	The user-assigned name for this PVC.

[1] The admin state or status under the **conf port iwf ds1/e1**, **conf fratm**, **conf fratm pvc**, **conf funi** and **conf funi pvc** levels reflects the configured state of a port, service, or connection. The admin status is unique to a level. For example, if a connection under **conf fratm pvc** is administered down, the corresponding admin state of the service under **conf fratm** is up, but the connection will not pass traffic. This is because the admin status of each of these entities has a direct effect on the operational status of all three entities. The operational status of these entities is currently not displayed in AMI.

You can also display information about a particular service as follows:

```
myswitch::configuration funi pvc> show ?
usage: show [<serviceId>] [<fvpi> [<fvci>]]] [advanced]
myswitch::configuration funi pvc> show 4a1:02
         FUNI FUNI Input        Output
 SvcId    VPI  VCI  Port VPI VCI  Port VPI VCI  Admin Eppd Name
 4A1:02    0   42   4A1  32  42   4C1   0  42    up    0  pvc_c
 4A1:02    0   42   4C1   0  42   4A1  32  42    up    0  pvc_c
```

The fields in this display are defined in the same manner as those in the previous example.

You can display information about a particular service, fvpi, and fvci as follows:

```
myswitch::configuration funi pvc> show 4a1:02 0 42
         FUNI FUNI Input        Output
 SvcId    VPI  VCI  Port VPI VCI  Port VPI VCI  Admin Eppd Name
 4A1:02    0   42   4A1  32  42   4C1   0  42    up    0  pvc_c
 4A1:02    0   42   4C1   0  42   4A1  32  42    up    0  pvc_c
```

The fields in this display are defined in the same manner as those in the previous example.

You can also display advanced information as follows:

```
myswitch::configuration funi pvc> show advanced
 SvcId  fvpi fvci afupc faupc
 4A1:00   0   40    0    0
 4A1:01   0   41    0    0
 4A1:02   0   42    0    0
```

The fields in this display are defined as follows：

Field	Description
SvcId	The ID of the service that is used by this PVC.
fvpi	The Virtual Path Identifier (VPI) of the FUNI PVC on the FUNI side.
fvci	The Virtual Channel Identifier (VCI) of the FUNI PVC on the FUNI side.
afupc	The UPC traffic contract that is used on this PVC in the ATM to Frame Relay direction.
faupc	The UPC traffic contract that is used on this PVC in the Frame Relay to ATM direction.

If no *FramePlus* network modules are installed or if no FUNI PVCs are configured, then the following is displayed:

```
myswitch::configuration funi pvc> show
No FUNI VCC information available
```

9.4   Configuring the Status of a FUNI Service

This command enables you to configure the administrative status of a particular FUNI service on a *FramePlus* network module to be up or down. Enter the following:

```
myswitch::configuration funi> admin <serviceid> (up|down)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The FUNI service whose administrative status needs to be changed. The ID number can be found in the SvcId field under **conf funi show**.
up | down	**up** changes the administrative status of the designated service to up (enabled). **down** changes the administrative status of the designated service to down (disabled).

9.5   Enabling and Disabling FUNI Service Statistics

This command allows you to enable or disable the collection of statistics on an existing FUNI service.

NOTE

To collect FUNI statistics, you must enable them both here <u>and</u> at the network module level under **conf module fram stats**. They are disabled here, but enabled at the module level by default.

Enter the following:

```
myswitch::configuration funi> stats <serviceid> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The FUNI service that is to be configured. The ID number can be found in the SvcId field under **conf funi show**.
enable | disable	**enable** means statistics are collected on the designated service. **disable** means statistics are not collected on the designated service. When statistics are disabled, the counters show zeros. The default is **disable**.

FUNI Configuration Commands

9.6   Enabling and Disabling FUNI Service Traps

This command allows you to to enable or disable the SNMP service UP and DOWN traps on an existing FUNI service. Enter the following:

```
myswitch::configuration funi> trap <serviceid> (enable|disable)
```

These parameters are defined as follows:

Parameter	Description
serviceid	The FUNI service that is to be configured. The ID number can be found in the `SvcId` field under **conf funi show**.
enable | disable	**enable** means SNMP UP and DOWN traps are operational on the designated service. **disable** means SNMP UP and DOWN traps are not operational on the designated service.

9.7   Displaying FUNI Service Information

To display information about the current FUNI services, enter the following:

```
myswitch::configuration funi> show
Searching for FUNI services
 SvcId  TimeSlot   Admin Funi  Serv Signal   Traps     Stats     Name
 4A1:00 1          up       1     0 nonexist enabled   disabled  service_a
 4A1:01 2          up       0     1 nonexist enabled   disabled  service_b
 4A1:02 3          up       0     0 nonexist enabled   disabled  service_c
```

The fields in this display are defined as follows：

Field	Description
SvcId	The ID for this service.
Timeslot	Indicates which timeslots (1-24 for DS1, 1-31 for E1) are being configured for this service. **all** indicates unchannelized connections. The time slot assignments must match the access rate specification in the service profile.
Admin[1]	The administrative status of this service. **up** means it is enabled. **down** means it is disabled.
Funi	The index of the FUNI profile that is to be used by this service.
Serv	The index of the service profile that is to be used by this service.
Signal	This is a read-only field that indicates whether or not there exists a FUNI signalling instance that is associated with this FUNI service.
Trap	**enabled** means SNMP traps are being generated for this service. **disabled** means SNMP traps are not being generated for this service.
Stats	**enabled** means statistics are being collected for this service. **disabled** means statistics are not being collected for this service.
Name	The user-assigned name for this service.

[1.] The admin state or status under the **conf port iwf ds1/e1**, **conf fratm**, **conf fratm pvc**, **conf funi** and **conf funi pvc** levels reflects the configured state of a port, service, or connection. The admin status is unique to a level. For example, if a connection under **conf fratm pvc** is administered down, the corresponding admin state under **conf fratm** is up, but the connection will not pass traffic. This is because the admin status of each of these entities has a direct effect on the operational status of all three entities. The operational status of these entities is currently not displayed in AMI.

If no *FramePlus* network modules are installed or if FUNI services have not been configured, then the following is displayed:

```
myswitch::configuration funi> show
Searching for FUNI services
....No services found
```

CHAPTER 10 HTTP Configuration Commands

These commands let you configure the location of the help files and enable or disable the server for the web-based Element Manager. Type **?** at the **http** level to display the available subcommands.

```
myswitch::configuration http> ?
   help_url          server          show
```

10.1 Changing the Location of the Help Files

By default, the help files for the web-based Element Manager are located on FORE Systems' external web site at `http://www.fore.com`. This command lets you specify a different location for the help files for the web-based Element Manager. Enter the following parameters:

```
myswitch::configuration http> help_url (default |<ip-address>[:port][<pathname>])
```

These parameters are defined as follows:

Parameter	Description
default	After you have specified a non-default location, entering **default** changes the location of the help files back to the default location of FORE Systems' external web site.
<ip-address>	The IP address of the web server to contact.
[:port]	The port on the web server to contact.
[<pathname>]	The directory on the web server into which the help files should be put.

Before changing the location of the help files, you first need to download the help files from FORE Systems' external web site, and place them on a host.

> **NOTE** ➤ The host that is to be used as a repository for the help files must be configured as a web server. Otherwise, the help files will not be accessible.

Once you have loaded the help files on the web server, you can use this command to change the location. For example, enter something similar to the following:

```
myswitch::configuration http> help_url 169.14.149.90/us/bob/helpfiles
```

If, at a later time, you want to change back to the default location, enter the following:

```
myswitch::configuration http> help_url default
```

10.2 Enabling/Disabling the HTTP Server

This command lets you enable or disable the HTTP server for the web-based Element Manager. When the server is enabled, more than one person can access and configure the switch at the same time.

> **NOTE** You need to have admin level privileges to execute this command. See Part 2 of this manual for more information about security privileges.

Enter the following parameters:

```
myswitch::configuration http> server (enabled | disabled)
```

These parameters are defined as follows:

Parameter	Description
enabled	**enabled** means the HTTP server is up and you can access the switch via the web-based Element Manager. This is the default.
disabled	**disabled** means the HTTP server is down and you cannot access the switch via the web-based Element Manager.

> **NOTE** If you change the password file, you should disable and enable the HTTP server, to force all current HTTP sessions to log in again.

10.3 Displaying the HTTP Server's Status

This command lets you display whether the HTTP server for the web-based Element Manager is enabled or disabled. Enter the following parameters:

```
myswitch::configuration http> show
HTTP Server is enabled
URL for Help Files is set to default
```

CHAPTER 11   ILMI SNMP Proxy Commands

These commands let you configure the ILMI SNMP Proxy (ISP) table. Through this table, you can discover the topology of the network to which your switch is connected because your switch registers its address via ILMI (if ILMI is running on each switch) with its neighboring switches. You can add, delete, or display ISP table entries and you can send SNMP requests to the ISP table. Type **ilmiproxy ?** at the **configuration** level to display the available sub-commands.

```
myswitch::configuration> ilmiproxy ?
  delete          go              new              show
```

11.1   Deleting an ISP Table Entry

This command lets you delete an entry from the ISP table. Enter the following parameters:

```
myswitch::configuration ilmiproxy> delete <port> <vpi> <index>
```

For example, to delete an entry from the table, enter something similar to the following:

```
myswitch::configuration ilmiproxy> del 1a2 0 19
```

11.2   Resending an ILMI SNMP Proxy Request

This command allows you to resend an ILMI SNMP proxy request. Enter the following parameters:

```
myswitch::configuration ilmiproxy> go <port> <vpi> <index>
```

For example, to resend a request, enter something similar to the following:

```
myswitch::configuration ilmiproxy> go 1a1 0 6
```

11.3   Creating an ISP Table Entry

This command lets you create an ISP table entry and send the request. Enter the following parameters:

```
myswitch::configuration ilmiproxy> new <port> <vpi> <index> (get | gnext) <oid> <comm>
(doOnce | doEveryIlmiRestart)
```

NOTE ▶ Your request will fail if ILMI is down on the path that you entered. First, use the command `conf signalling show` so that you can quickly see on which ports ILMI is up.

For example, before creating an entry, first check on which ports ILMI is up by entering the following:

```
myswitch::configuration signalling> show
 Port VPI Interface       SigVersion State ILMI Side    RemoteAddress
 1C1    0 privateUNI(a)  uni30(a)   down  down network
 1C2    0 privateUNI     uni31      up    up   network 172.19.12.140
 1C3    0 PNNI(a)        pnni10(a)  up    up   network 172.19.12.57
 1C4    0 PNNI           pnni10     up    up   network 169.144.64.58
 1CTL   0 privateUNI(a)  uni30(a)   up    down network
```

Then type the parameters to create the entry as follows:

```
myswitch::configuration ilmiproxy> new 1c2 0 get .1.3.6.1.2.1.1.4.0 public doOnce
```

The parameters for delete, go, and new are defined as follows:

Parameter	Description
port	The port number of the ISP table entry.
vpi	The VPI of the ISP table entry.
index	The unique index number of the ISP table entry.
get | getnext	`get` means the SNMP request you are creating is a get, which searches for the object that you are requesting. `gnext` means the SNMP request you are creating is a getnext, which searches for the next greatest entry in the MIB after the object that you specify.
oid	The object identifier (OID) for the SNMP request that you want to get or getnext.
comm	The SNMP community string to use.
doOnce	The SNMP request is to be performed only once.
doEveryIlmiRestart	The SNMP request is to be performed each time that ILMI restarts on the specified path.

11.4   Displaying the ISP Table Entries

This command lets you display the current ISP table entries. Enter the following parameters:

```
myswitch::configuration ilmiproxy> show
ISP Entry Port 1A1 Vpi 0 Index 3
-----------------------------------------------------------------------
Req_Oid          .1.3.6.1.2.1.1.2.0
Rsp_Oid          .1.3.6.1.2.1.1.2.0
Operation        get
Community        private
When to do       doOnce
Value            OID:.1.3.6.1.4.1.326.2.2
Operation Status success
-----------------------------------------------------------------------
ISP Entry Port 1A1 Vpi 0 Index 12
-----------------------------------------------------------------------
Req_Oid          .1.3.6.1.4.1.326.2.1.1.1.1.16.0
Rsp_Oid          N/A
Operation        get
Community        private
When to do       doOnce
Value            N/A
Operation Status failure
-----------------------------------------------------------------------
ISP Entry Port 1A1 Vpi 0 Index 15
-----------------------------------------------------------------------
Req_Oid          .1.3.6.1.2.1.1.1.0
Rsp_Oid          .1.3.6.1.2.1.1.1.0
Operation        get
Community        public
When to do       doOnce
Value            FORE Systems ASX-200BX
Operation Status success
-----------------------------------------------------------------------
Press return for more, q to quit: q
```

The fields in this display are defined as follows：

Field	Description
Req_Oid	The object identifier (OID) for the SNMP request that you sent.
Rsp_Oid	The OID in the response to the SNMP request. It is valid only if the operation status is `success`.
Operation	**get** means the SNMP request is a get, which searches for the OID that you requested. **gnext** means the SNMP request is a getnext, which searches for the next greatest entry in the MIB after the specified OID.
Community	The SNMP community string being used. The default is public.
When to do	`doOnce` means that the SNMP request is performed only once. `doEveryIlmirestart` means that the SNMP request is to be performed each time that ILMI comes up on the specified path.
Value	The value returned by the get or the getnext. This is valid only if the operation status is `success`
Operation Status	The current status of this SNMP query. **Idle** means the request has not been made yet. `inProgress` means the request has been made, but has not been completed yet. **Success** means the request has been made and completed successfully. **Failure** means the request has been made, but was not successful. Failures can occur because either there was no response from the peer (the request timed out), or there was an error message from the peer (like a NOSUCHNAME error), or the ILMI/UNI/link on which to send this request is down.

You can also display an individual ISP table entry. Enter the following:

```
myswitch::configuration ilmiproxy> show [<port> [<vpi> [<index>]]]
myswitch::configuration ilmiproxy> show 1a1 0 15
ISP Entry Port 1A1 Vpi 0 Index 15
---------------------------------------------------------------------
Req_Oid          .1.3.6.1.2.1.1.1.0
Rsp_Oid          .1.3.6.1.2.1.1.1.0
Operation        get
Community        public
When to do       doOnce
Value            FORE Systems ASX-200BXTNX-210
Operation Status success
---------------------------------------------------------------------
```

If there are no current entries in the ISP table, then the following is displayed:

```
myswitch::configuration ilmiproxy> show
ISP information not available
```

CHAPTER 12   IP Configuration Commands

These commands let you change the IP configuration. You can display the list of available sub-commands by typing **ip ?** at the **configuration** level.

```
myswitch::configuration> ip ?
  address           admin           broadcast          forwarding
  mtu               route>          show               unconfigure
```

12.1   Configuring the IP Address

This command allows you to configure the IP address, netmask, and administrative status of each of the switch's IP interfaces. Enter the following parameters:

NOTE On a new switch, the `ie0`, `asx0`, `qaa0`, `qaa1`, `qaa2`, `qaa3` interfaces are NOT configured. An IP address must be configured for at least one of the interfaces to allow IP access to the switch, which, in turn, enables SNMP access. By setting the IP address of the `asx0` interface or one of the `qaa` interfaces, in-band (over ATM) access to the switch control processor (SCP) is enabled. By setting the IP address of the `ie0` interface, out-of-band access to the SCP is enabled.

NOTE On an ASX-1000 or a TNX-1100, the IP addresses must be configured individually on each SCP.

```
myswitch::configuration ip> address <interface> <address> [<netmask>] [(up)]
```

These parameters are defined as follows:

Parameter	Description
interface	The name of the IP interface to be managed. Valid interfaces are: **ie0** (the Ethernet interface), **asx0** (the switch's SPANS interface), **qaa0**, **qaa1**, **qaa2**, **qaa3** (the Classical IP interfaces), **lo0** (the switch's local interface that allows AMI to run), and **el0**, **el1**,etc. (the LAN Emulation interfaces).
address	The IP address for this interface. The state of the interface must be **up** before setting the address. This can be changed using **conf ip admin**.
netmask	The subnet mask for this IP interface. It should be entered in dotted decimal notation (e.g., 255.255.255.0). If you accidentally enter a duplicate address for the same subnet as an existing one, the switch rejects the duplicate and prints an error message. If you do not provide a netmask, then a default netmask is used.
up	Entering **up** brings the designated interface on-line as soon as you enter this command.

12.2   Configuring the IP State

This command allows you to change the state of the IP interface to up or down. Enter the following parameters:

```
myswitch::configuration ip> admin <interface> (up|down)
```

These parameters are defined as follows:

Parameter	Description
interface[1]	The name of the IP interface to be managed. Valid interfaces are: **ie0** and **asx0**.
address	**up** brings the designated interface on-line. **down** takes the interface off-line. If you enter **conf ip admin el all down**, you can bring down all of the **el** interfaces at once. Similarly, entering **conf ip admin el all up** brings up all of the **el** interfaces at once.

[1.] The switch's local interface, lo0, must <u>always</u> be up to allow AMI to run on the switch.

12.3   Configuring the IP Broadcast Address

This command lets you modify the broadcast address for an IP interface. Enter the following:

```
myswitch::configuration ip> broadcast <interface> (0|1)
```

These parameters are defined as follows:

Parameter	Description
interface	The name of the IP interface to be managed. Valid interfaces are: **ie0**, **asx0**, and any of the **el** interfaces.
0 | 1	The IP broadcast type for this interface. This is the host portion of the IP address that is used for routing. **1** causes the host portion of the IP address to be set to all 1s. **0** causes the host portion of the IP address to be set to all 0s.

IP Configuration Commands

12.4   Configuring IP Forwarding

This command allows you to turn IP forwarding on or off. If IP forwarding is turned off, the switch will not forward (i.e., route) IP packets from one IP interface to another IP interface. Enter the following parameters:

```
myswitch::configuration ip> forwarding (on|off)
```

These parameters are defined as follows:

Parameter	Description
on | off	Using `on` turns IP forwarding on. Using `off` turns IP forwarding off. The default is `off`.

12.5   Configuring the MTU Size of a Classical IP Interface

This command allows you to configure the MTU size of a Classical IP interface. Although the default MTU size should work in most installations, you may need to modify this value when connected to non-FORE equipment. Enter the following parameters:

```
myswitch::configuration ip> mtu <qaaX> <size>
```

These parameters are defined as follows:

Parameter	Description
qaaX	The name of the Classical IP interface. Valid interfaces are: `qaa0`, `qaa1`, `qaa2`, and `qaa3`.
size	The MTU size for this Classical IP interface. Valid values are from 1 to 32767. The default is `9180`.

12.6   Configuring IP Routes

This command allows you to add a static IP route to the local IP routing table, delete a static IP route from the local IP routing table, or list the current static IP routes in the local IP routing table. You can display the list of available subcommands by typing **route ?** at the **ip** level.

```
myswitch::configuration ip> route ?
new                delete             show
```

12.6.1 Adding an IP Route

This command lets you create an IP route. Enter the following parameters:

```
myswitch::configuration ip route> new (default | <destination-ipaddress>) <gateway>
[<metric>] [(host | net)]
```

12.6.2 Deleting an IP Route

This command lets you delete an IP route. Enter the following parameters:

```
myswitch::configuration ip route> delete (default | <destination-ipaddress>) <gateway>
```

These parameters for new and delete are defined as follows:

Parameter	Description
default	This parameter must be entered to create a default route.
destination-ipaddress	The destination IP network number.
gateway	The gateway address to the destination IP network number.
metric	The number of hops to the destination IP network. If 1 is specified, the route is created with the RTF_GATEWAY flag set. The default is 1.
host | net	**host** means this is a host-specific route with the destination being a specific node's IP address. **net** means this is a network-specific route with the destination being a network IP address. The default is **net**.

12.6.3 Showing the IP Routes

This command lets you display the current IP routes. Enter the following parameters:

```
myswitch::configuration ip route> show
Destination     Gateway         Metric      Interface       Flags
default         169.144.48.1    1           le0             G
169.144.48.0    169.144.48.21   0           le0
169.144.60.0    169.144.60.21   0           asx0
169.144.64.0    169.144.64.21   0           qaa0
169.144.200.0   169.144.200.21  0           el0
169.144.204.0   169.144.204.21  0           el1
169.144.205.0   169.144.205.21  0           el2
169.144.206.0   169.144.206.21  0           el3
```

The fields in this display are defined as follows:

Field	Description
Destination	The destination IP network.
Gateway	The gateway address to the destination IP network number.
Metric	The number of hops to the destination IP network. The default is 1.
Interface	The local IP interface used to get to the destination IP network.
Flags	H means the route is host-specific (created with the RTF_HOST flag set). G means the route is network-specific (created with the RTF_GATEWAY flag set).

12.7   Displaying the IP Interface Configuration

This command allows you to display information about the configuration of the IP interfaces. Enter the following parameters:

```
myswitch::configuration ip> show
interface  state       address        netmask        broadcast        mtu
lo0        up          127.0.0.1      255.0.0.0      N/A              4096
ie0        up          169.144.229.45 255.255.255.0  169.144.229.255  1500
asx0       down
qaa0       up          169.144.230.45 255.255.255.0  N/A              9180
qaa1       down
qaa2       down
qaa3       down

IP Forwarding State: not-forwarding
```

The fields in this display are defined as follows:

Field	Description
interface	The name of the IP interface.
state	The administrative state of the IP interface.
address	The IP address of the IP interface.
netmask	The netmask address of the IP interface.
broadcast	The broadcast address of the IP interface.
mtu	The MTU size of the IP interface. Only the Classical IP interfaces (qaa0, qaa1, qaa2, and qaa3) have a configurable MTU size.

You may also designate a single interface to be displayed by entering **show** and the specific interface name at the prompt as follows:

```
myswitch::configuration ip> show ie0
interface  state       address        netmask        broadcast        mtu
ie0        up          169.144.229.45 255.255.255.0  169.144.229.255  1500

IP Forwarding State: not-forwarding
```

The fields in this display are defined in the same manner as those listed above in the example for showing the configuration of all of the IP interfaces.

12.8   Unconfiguring an IP Interface

This command lets you unconfigure an IP interface. This process removes the IP address associated with the interface. Enter the following parameters:

```
myswitch::configuration ip> unconfigure <interface>
```

This parameter is defined as follows:

Parameter	Description
interface	The name of the IP interface to be unconfigured. Valid interfaces are: **ie0**, **asx0**, **qaa0**, **qaa1**, **qaa2**, and **qaa3**.

You will be asked to confirm this action. To confirm the action, type **y** at the prompt. To cancel the action, type **n** or press **<Enter>** at the prompt. For example:

```
myswitch::configuration ip> unconfigure qaa1

Unconfiguring an interface requires the switch to be rebooted.
Continue with unconfigure [n]? y

Reboot the switch [y]? y
```

NOTE The switch must be rebooted for this command to take effect.

CHAPTER 13 LANE Configuration Commands

These commands allow you to configure LAN Emulation (LANE) on a switch. You can display the list of available subcommands by typing **?** at the **lane** level.

```
myswitch::configuration lane> ?
  default>         bus>              lec>              lecs>
  les>
```

13.1   Default LANE Configuration Commands

These commands let you configure a default ELAN. You can display the list of available subcommands by typing **?** at the **default** level.

```
myswitch::configuration lane default> ?
  new              delete          show
```

13.1.1 Creating a Default ELAN

This command lets you create a simple default ELAN. You can use this command only if you do not already have an ELAN named `default`. This command creates and starts an ELAN named `default` that consists of an LECS, a co-located LES/BUS, and a single LEC. This ELAN uses the ATM Forum well-known address for the LECS. It is an Ethernet ELAN with an MTU size of 1516, and it has ELAN access control disabled and TLV registration enabled.

> **NOTE** You must create a LECS configuration file using **conf lane lecs new** and put it on the switch that is to be the LECS using **conf lane lecs get** *<host>:<remotefile>* [*<localfile>*].

> **NOTE** You must assign an IP address to the `el0` interface for the LEC using **conf ip address** and make the LEC operational by using **conf ip admin** *<interface>* **up**.

> **NOTE** You cannot use DLE with this ELAN unless you delete the ELAN and recreate it using the **conf lane les** command so that you can assign the DLE peer server addresses.

Enter the following parameters:

```
myswitch::configuration lane default> new
Do you want to create default ELAN (LES/BUS, LECS, LEC): [n]? y
```

Entering **n** or pressing **<ENTER>** aborts the command. Entering **y** creates a new ELAN named `default`. If you enter **y**, you receive the following message when the ELAN is created.

```
Created LANE Services and a LEC for the default ELAN.
```

If an ELAN named `default` already exists, you will receive an error message.

13.1.2 Deleting a Default ELAN

This command lets you delete the ELAN named default. This command will delete the ELAN named default that was started using the **conf lane default new** command.

> **NOTE** If the ELAN named default was started using the other **conf lane** commands and the -les option was specified under **conf lane lecs new**, then this command will delete that ELAN named default, as well.

Enter the following parameters:

```
myswitch::configuration lane default> delete
Do you really want to delete the ELAN named default (LES/BUS, LECS, LEC): [n]? y
```

Entering **n** or pressing **<ENTER>** aborts the command. Entering **y** deletes the ELAN named default. If no ELAN named default exists, you receive the following message when you enter **y**.

```
myswitch::configuration lane default> delete
Do you really want to delete the ELAN named default (LES/BUS, LECS, LEC): [n]? y
Default LANE is not configured.
```

13.1.3 Displaying a Default ELAN

This command lets you show information for the ELAN named `default`. Enter the following:

```
myswitch::configuration lane default> show
LECS Information:
 Index  AdminStatus  OperStatus   Selector  WKA         Database
     1  up           up                0xf1 atm-forum   lecs.cfg
         Default LES :0x47.0005.80.ffe100.0000.f21a.3596.0020481a3596.f0
LES/BUS Information:
 Index AdminStatus OperStatus LesSel Type       MTU   ELAN      SECURE  TLVs
 1     up          up            0xf0 ethernet   1516  default   disable enable
         LES : 0x47.0005.80.ffe100.0000.f21a.3596.0020481a3596.f0
         BUS : 0x47.0005.80.ffe100.0000.f21a.3596.0020481a3596.f0 (Co-Located)
LEC Information:
         Admin   Oper
 Index  Status  Status  Sel  Mode        MACaddress   IfName   ELAN
     1  up      up       0x00 wellknown   0620481a3596 el0      default
         LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
         LES :0x47.0005.80.ffe100.0000.f21a.3596.0020481a3596.f0
```

The fields in this display are defined as follows:

Field	Description
LECS Information:	
Index	The unique index number of the LECS that identifies it from other LECSs. This number is dynamically assigned by the switch when the LECS is created.
AdminStatus	Reflects any changes that you have made to the status of the LECS. **up** means you have started the LECS. **down** means you have stopped the LECS.
OperStatus	Reflects the actual current status of the LECS. **up** means the LECS is currently active. **down** means the LECS is currently inactive.
Selector	The selector byte portion (20th byte) of the ATM address, in hexadecimal format, of the switch that is running the LECS.
WKA	The state of the well-known address for this particular LECS. **atm-forum** means the LECS is using the default address as defined by the ATM Forum. This is the default value. **none** means the well-known address has been disabled. **other** means the well-known address has been redefined, in which case the new address appears on the following line.
Database	The full path to the location and name of the LECS database file. The default is `lecs.cfg`.
Default LES	The default LES address to use in case the LECS configuration file is inaccessible. If a default LES address has not been specified, this field is not displayed.
LES/BUS Information:	
Index	The unique index number of the LES that identifies it from other LESs. This number is dynamically assigned by the switch when the LES is created.

Field	Description
AdminStatus	Reflects any changes that you have made to the status of the LES. **up** means you have started the LES. **down** means you have stopped the LES.
OperStatus	Reflects the actual current status of the LES. **up** means the LES is currently active. **down** means the LES is currently inactive.
LesSel	The selector byte portion (20th byte) of the ATM address, in hexadecimal format, of the switch that is the LES.
Type	**ethernet** means the LAN type is Ethernet. **token-ring** means the LAN type is Token Ring.
MTU	The maximum data frame size. Can be **1516**, **1580**, **4544**, **9234**, or **18190**. The default for Ethernet is **1516**. The default for Token Ring is **4544**.
ELAN	The name of the ELAN that this LES services.
SECURE	**enable** means that ELAN access control is running. **disable** means that ELAN access control is not running.
TLVs	**enable** means the LES accepts TLV parameters from LEC/MPCs that register with it and distributes these TLVs to LEC/MPCs in response to LE-ARP queries. **disable** means the LES collects TLV parameters from LEC/MPCs that register with it, but does not distribute these TLVs to LEC/MPCs in response to LE-ARP queries.
LES:	The full ATM address of the LES. The LECs that wish to join this ELAN use this address to contact the LES.
BUS:	The full ATM address of the BUS for this ELAN.
LEC Information:	
Index	The unique index that identifies this LEC. It is dynamically assigned by the switch when the LEC is created.
AdminStatus	Reflects any changes that you have made to the status of the LEC. **up** means you have started the LEC. **down** means you have stopped the LEC.
OperStatus	Reflects the actual current status of the LEC. **up** means the LEC is currently active. **down** means the LEC is currently inactive. **joining** means that the LEC is in the process of registering with the ELAN.
Sel	The selector byte portion (20th byte) of the ATM address of the LEC in hexadecimal format.
Mode	The configuration mode that is used when a LEC joins the ELAN. **wellknown** means that the "well-known" LECS address and the default LES are used. **manual** means that the specified LECS or LES address is used. The default is **wellknown**.
MACaddress	The Ethernet MAC address for this LEC.
IfName	The **el** interface name of this LEC.
ELAN	The name of the ELAN to which this LEC belongs.

If there is no ELAN named `default`, you receive the following message:

```
myswitch::configuration lane default> show
Default LANE is not configured.
```

LANE Configuration Commands

13.2   BUS Configuration Commands

These commands allow you to configure a Broadcast and Unknown Server (BUS) for an ELAN. You can display the list of available subcommands by typing **?** at the **bus** level.

```
myswitch::configuration lane bus> ?
  admin           delete          new             show
```

NOTE ➤ *ForeThought* 5.0 and greater versions do not allow you to create a BUS separately from a LES. Therefore, the commands in this menu are only useful in providing backwards compatibility with switches that are running earlier versions of *ForeThought* software.

13.2.1 Configuring the BUS Administrative Status

This command lets you change the administrative status of a BUS to up (start a BUS service) or down (stop a BUS service). Enter the following parameters:

```
myswitch::configuration lane bus> admin <BUS index | BUS index range (x-y)> (up | down)
```

13.2.2 Deleting a BUS

This command allows you to delete a specified BUS. Enter the following parameters:

```
myswitch::configuration lane bus> delete <BUS index | BUS index range (x-y)>
```

The parameters for admin and delete are defined as follows:

Parameter	Description
BUS index	The unique, positive integer index of the BUS that is dynamically assigned by AMI when a BUS is created to identify this service from any other service in the same class. Found under the Index field using the **conf lane bus show** command.
BUS index range (x-y)	The range of index numbers of the BUSs that you want to start or stop. The index is found under the Index field using **conf lane bus show**.
up|down	**up** changes the administrative status of the designated BUS index to up. **down** changes the administrative status of the designated BUS index to down.

13.2.3 Creating a BUS

This command lets you create a BUS for an ELAN. Enter the following parameters:

```
myswitch::configuration lane bus> new <BUS Selector byte (HEX)> <BUS name> \
              [-type (ethernet | token-ring)] \
              [-mtu (1516 | 4544 | 9234 | 18190)]
```

The parameters for new are defined as follows:

Parameter	Description
BUS Selector byte	The 20th byte of the ATM address of the switch that is to run a BUS service (entered in hexadecimal format). Use **conf atmarp getnsap** to display the entire ATM address.
BUS name	The name for this BUS. Identifies which ELAN this BUS services.
type	**ethernet** means that the LAN type is Ethernet. **token-ring** means that the LAN type is Token Ring. The default is ethernet.
mtu[1]	Indicates which maximum transmission unit (MTU) size you wish to use. Valid values are: 1516, 4544, 9234, and 18190. Defaults are 1516 for Ethernet and 4544 for Token Ring.

[1.] The MTU size must match the MTU size of the other hosts and edge devices on the ELAN.

13.2.4 Displaying BUS Information

This command lets you display the current BUS information. To display information about every BUS that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane bus> show
  Index  AdminStatus  OperStatus    Selector  Type        MTU   ELAN
      1  up           up               0x12   ethernet    1516  one
      2  up           up               0x32   token-ring  4544  three
```

The fields in this display are defined as follows:

Field	Description
Index	The unique index number that identifies this BUS. It is dynamically assigned by AMI when the BUS is created.
AdminStatus	Reflects any changes that you have made to the status of the BUS. **Up** means you have started the BUS. **Down** means you have stopped the BUS.
OperStatus[1]	Reflects the actual current status of the BUS. **Up** means the BUS is currently active. **Down** means the BUS is currently inactive.
Selector	The selector byte portion (20th byte) of the ATM address of the host or switch that is the BUS in hexadecimal format
Type	The type of ELAN this is. **ethernet** means that the LAN type is Ethernet. **token-ring** means that the LAN type is Token Ring
MTU	The maximum transmission unit (MTU) size.
ELAN	The name of the ELAN that this BUS services.

[1.] When you change the administrative status of a BUS from down to up, it takes a few seconds for the change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the information above to show the AdminStatus as up, but the OperStatus as down. If you refresh the display, the change will have taken place and be reflected here.

To display information about a particular BUS that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane bus> show [<BUS index>]
```

For example, to display information about the BUS with an index number of 1, enter the following parameters:

```
myswitch::configuration lane bus> show 1
  Index  AdminStatus  OperStatus    Selector  Type        MTU   ELAN
      1  up           up               0x12   ethernet    1516  one
```

The fields in this display are defined in the same manner as those in the previous example.

13.3 LEC Configuration Commands

These commands let you configure the LAN Emulation Client (LEC). You can display the list of available subcommands by typing **lec ?** at the **lane** level.

```
myswitch::configuration lane> lec ?
  admin           arp>            delete          default>
  new             show
```

13.3.1 Configuring the LEC Administrative Status

This command lets you change the administrative status of a LEC to up (start a LEC) or down (stop a LEC). Enter the following parameters:

```
myswitch::configuration lane lec> admin <LEC index | LEC index range (x-y)> (up | down)
```

The parameters for admin are defined as follows:

Parameter	Description
LEC index	The unique, positive integer index of the LEC that is dynamically assigned by AMI when a LEC is created to identify this LEC from any others in the same ELAN. Found under the Index field using the **conf lane lec show** command.
LEC index range (x-y)	The range of index numbers of the LECs that you want to start or stop. The index is found under the Index field using **conf lane lec show**.
up | down	Entering **up** starts this LEC. Entering **down** stops this existing LEC. The default is **up**.

13.3.2 Configuring LANE ARP Commands

These commands let you configure the LANE ARP cache. You can reach this level by entering **arp** at the **lec** level. Enter the following parameters to list the various ARP commands:

```
myswitch::configuration lane lec> arp ?
  delete          show
```

13.3.2.1 Deleting LANE ARP Cache Information

This command allows you to remove an ARP entry from the LANE ARP cache or to delete the contents of the LANE ARP cache. Enter the following parameters:

```
myswitch::configuration lane lec arp> delete (all | <MAC address>)
```

These parameters are defined as follows:

Parameter	Description
all	Indicates that all of the entries are to be flushed from the LANE ARP cache.
MAC address	Indicates the specific entry that is to be flushed from the LANE ARP cache.

13.3.2.2 Displaying LANE ARP Cache Information

This command displays the current LANE ARP cache. The MAC address-to-ATM address mapping information for each LEC is stored here. Enter the following parameters:

```
myswitch::configuration lane lec arp> show [(advanced)]
```

By entering **show** without the **advanced** option, the basic LANE ARP cache information is displayed as follows:

```
myswitch::configuration lane lec arp> show
 MacAddress    AtmAddress                                         ELAN
  0020481a00d5 0x47.0005.80.ffe100.0000.f21a.00d5.0020481a00d5.0b eng_net
```

By entering **show** with the **advanced** option, more LANE ARP cache information, including the VPI/VCI combination and any flags associated with this entry, is displayed as follows:

```
myswitch::configuration lane lec arp> show advanced
 MacAddress    AtmAddress                                         ELAN
  0020481a00d5 0x47.0005.80.ffe100.0000.f21a.00d5.0020481a00d5.0b eng_net
          vpi=0, vci=82, flags=valid
```

If the LANE ARP cache is empty, then the following message is displayed.

```
No LANE ARP entries are available.
```

13.3.3 Deleting a LEC

This command lets you delete a LEC from an ELAN. Enter the following parameters:

```
myswitch::configuration lane lec> delete <LEC index | LEC index range (x-y)>
```

The parameters for delete are defined as follows:

Parameter	Description
LEC index	The unique, positive integer index of the LEC that is dynamically assigned by AMI when a LEC is created to identify this LEC from any others in the same ELAN. The index is found under the Index field using the **conf lane lec show** command.
LEC index range (x-y)	The range of index numbers of the LECs that you want to delete. The index is found under the Index field using **conf lane lec show**.

> **NOTE**
>
> Before you can delete a LEC, you must first administer its el interface down using the **conf ip admin** <*interface*> **down** command.

13.3.4 LEC Default Configuration Mode Commands

These commands allow you to set or to display the default LEC configuration mode. You can show the list of available subcommands by typing **default ?** at the **lec** level.

```
myswitch::configuration lane lec> default ?
  mode              show
```

13.3.4.1 Setting the Default LEC Configuration Mode

This command lets you set the default mode for configuring all of the ELANs that may be created on this switch.

> **NOTE**
>
> If you use **manual** mode, you <u>must</u> specify the LECS address of the machine that will be used as the LECS. If you use **wellknown** mode, then the "well-known" LECS address is used.

Enter the following parameters:

```
myswitch::configuration lane lec default> mode (manual | wellknown) [<LECS address>]
        LECS address is required for manual mode.
```

These parameters are defined as follows:

Parameter	Description
manual | wellknown	**manual** means the LECS address specified here is used as the LECS address. **wellknown** means that the LEC first uses ILMI to attempt to discover the LECS address. If no address is available via that method, then the LEC attempts to use the ATM Forum's "well-known" LECS address (either **47.0079.00.000000.0000.0000.0000.00A03E000001.00** or **c5.0079.00.000000.0000.0000.0000.00A03E000001.00**) to contact the LECS. The default is **wellknown**.
LECS address	The ATM address of the LECS to be used instead of the "well-known" LECS.

13.3.4.2 Displaying the Default LEC Configuration Mode

This command lets you show whether the default LEC configuration mode is **manual** (using a LECS other than the one at the "well-known" address) or **wellknown** (using the LECS at the "well-known" address). Enter the following parameters:

```
myswitch::configuration lane lec default> show
LEC Default configuration mode: wellknown
```

13.3.5 Creating an Ethernet LEC

This command lets you create an Ethernet LEC (join an ELAN). When a LEC is created, a corresponding `el` interface is created. The interface name (`el0`, `el1`, etc.) is assigned based on the selector byte entered when the LEC is created. The list of current `el` interfaces can be displayed using the **conf lane lec show** command or the **conf ip show** command. The maximum number of LECs that can be created on a switch is 16. Enter the following:

> **NOTE** ▶ This command only allows you to create an instance of a LEC on a switch. To create a LEC on a host, you must use the *ForeRunner* VLAN Manager or use an *ForeRunner* adapter. Refer to the respective user's manual for more information.

> **NOTE** ▶ You can only create an Ethernet LEC on a switch. To create a Token Ring LEC, you must use a *ForeRunner* PC adapter. Refer to your respective PC adapter user's manual for more information.

```
myswitch::configuration lane lec> new <LEC Selector byte (HEX)> <ELAN name>
        [-ip <IP Address> [-mask <IP netmask>]]
        [(wellknown | manual)]
        manual mode options: [-lecs <LECS address>] or [-les <LES address>]

        LES address should be an anycast address when using DLE.
        Use ELAN name "default" to join the default ELAN
```

The parameters for new are defined as follows:

Parameter	Description
LEC Selector byte (HEX)	The 20th byte of the ATM address of the LEC (entered in hexadecimal format). Use the **conf atmarp getnsap** command to display the entire ATM address.
ELAN name	The name of the ELAN that this LEC is joining.
ip <IP Address>	The IP address for this el interface. By entering this information here, you will not have to bring up the LEC's el interface separately using **conf ip admin** <*interface*> **up**.
mask <IP netmask>	The optional subnet mask for this el interface that shows the significant part of the address. It should be entered in dotted decimal notation (e.g., 255.255.255.0).

Parameter	Description
wellknown | manual	The configuration mode that is used when this LEC joins the ELAN.
	Using `wellknown` means that the LEC first attempts to discover the LECS address using ILMI. If the LEC discovers more than one LECS address, the LEC tries to connect to the LECS in the order that the addresses appear in the list. (This list can be configured using `conf nsap registry new` and displayed using `conf nsap registry show`.) If there are no addresses in this list, then the LEC attempts to contact the LECS using the ATM Forum "well-known" address.
	Using `manual` and the `-lecs` option means that the LEC attempts to contact the LECS using the LECS address you specified here.
	Using `manual` and the `-les` option means that the LEC bypasses the LECS and directly contacts the LES address specified here. If you are using DLE, then this LES address must be the anycast address.
	The default is `wellknown`.
-lecs <LECS address>	The LECS address to use instead of the "well-known" LECS address.
-les <LES address>	The LES address to use for this ELAN. If you are using DLE, then this LES address must be the anycast address.

NOTE If you create a LEC and then change the NSAP prefix of the switch at a later time, the change is NOT automatically reflected in the LEC's address. You must delete the LEC and recreate it using the new NSAP prefix.

The following is an example of how to create a LEC:

```
myswitch::configuration lane lec> new 0x0 eng_net -ip 169.144.224.222 -mask
255.255.255.0
```

If you decide not to enter the IP address at the time that you create the LEC, the switch will remind you that you still need to do this. For example:

```
myswitch::configuration lane lec> new 0x0 eng_net

  LEC is created. Use the commands in "configuration ip" menu
  to configure this interface.
```

13.3.6 Displaying LEC Information

This command lets you display the current LEC information. To display information about all of the LECs that are currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane lec> show
        Admin   Oper
 Index  Status  Status   Sel   Mode        MACaddress     IfName    ELAN
     1  up      up       0x00  wellknown   0220481c10bb   el0       eng_net
        LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
        LES :c5.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.20
```

The fields in this display are defined as follows:

Field	Description
Index	The unique index that identifies this LEC. It is dynamically assigned by the switch when the LEC is created.
Admin Status	Reflects any changes that you have made to the status of the LEC. `Up` means you have started the LEC. `Down` means you have stopped the LEC.
Oper Status[1]	Reflects the actual current status of the LEC. `Up` means the LEC is currently active. `Down` means the LEC is currently inactive. `Joining` means that the LEC is in the process of registering with the ELAN.
Sel	The selector byte portion (20th byte) of the ATM address of the LEC in hexadecimal format.
Mode	The configuration mode that is used when a LEC joins the ELAN. **wellknown** means that the LECS address was discovered via ILMI or that ATM Forum well-known address is being used. **manual** means that the specified LECS or LES address is used. The default is **wellknown.**
MACaddress	The Ethernet MAC address for this LEC.
IfName	The `el` interface name of this LEC.
ELAN	The name of the ELAN to which this LEC belongs.
LECS	The address of the LECS that services the ELAN to which this LEC belongs. If the `Mode` is **manual** then no LECS address is displayed because the LEC directly connects to the LES. If the `Mode` is **wellknown** then the LECS address is displayed. If the LECS address was discovered via ILMI, then that LECS address is displayed; otherwise, the ATM Forum well-known address is displayed.
LES	The address of the LES that services the ELAN to which this LEC belongs. If you are using DLE, this is the anycast address.

[1.] When you change the administrative status from `down` to `up`, it takes a few seconds for the change to occur and to be reflected in the `OperStatus` field. Therefore, it is possible for the `AdminStatus` to be `up`, but the `OperStatus` to be `down`. If you refresh the display, the change has taken place and is reflected here.

To display information about a particular LEC that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane lec> show [<LEC index>] [(advanced)]
```

For example, to display information about the LEC with an index number of 1, enter the following parameters:

```
myswitch::configuration lane lec> show 1
        Admin   Oper
 Index  Status  Status  Sel   Mode      MACaddress   IfName   ELAN
     1  up      up       0x00  wellknown 0220481c10bb  el0      eng_net
        LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
        LES :0xc5.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.20
```

The fields in this display are defined in the same manner as those in the previous example.

You can also display advanced information about the LECs. Enter the following parameters:

```
myswitch::configuration lane lec> show advanced
        Admin   Oper
 Index  Status  Status  Sel   Mode      MACaddress   IfName   ELAN
     1  up      up       0x00  wellknown 0220481c10bb  el0      eng_net
        LECS:0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
        LES :0xc5.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.20
        BUS : 47000580ffe1000000f21c10bb0020481c10bb20
        LEC ID : 2                         Discovered ELAN name : eng_net
        Configure Direct VCC : 0.71        Maximum Frame Size : 1516
        Control Direct VCC : 0.125         Control Distribute VCC : 0.75
        Multicast Send VCC : 0.128         Multicast Forward VCC : 0.76
        Last Error : ran out of resources (to LES)
```

The fields in this display are defined as follows:

Field	Description
Index	The unique index that identifies this LEC. It is dynamically assigned by the switch when the LEC is created.
Admin Status	Reflects any changes that you have made to the status of the LEC. Up means you have started the LEC. Down means you have stopped the LEC.

Field	Description
Oper Status[1]	Reflects the actual current status of the LEC. `Up` means the LEC is currently active. `Down` means the LEC is currently inactive. `Joining` means that the LEC is in the process of registering with the ELAN.
Sel	The selector byte portion (20th byte) of the ATM address of the LEC in hexadecimal format.
Mode	The configuration mode that is used when a LEC joins the ELAN. **wellknown** means that the LECS address was discovered via ILMI or that ATM Forum well-known address is being used. **manual** means that the specified LECS or LES address is used. The default is **wellknown**.
MACaddress	The Ethernet MAC address for this LEC.
IfName	The `el` interface name of this LEC.
ELAN	The name of the ELAN to which this LEC belongs.
LECS	The address of the LECS that services the ELAN to which this LEC belongs. If the `Mode` is **manual** then no LECS address is displayed because the LEC directly connects to the LES. If the `Mode` is **wellknown** then the LECS address is displayed. If the LECS address was discovered via ILMI, then that LECS address is displayed; otherwise, the ATM Forum well-known address is displayed.
LES	The address of the LES that services the ELAN to which this LEC belongs. If you are using DLE, this is the anycast address.
BUS	The address of the BUS that services the ELAN to which this LEC belongs.
LEC ID	The unique ID that the LES gives the LEC when it joins the ELAN.
Discovered ELAN Name	The ELAN name returned to the LEC by the LES when it joins the ELAN.
Configure Direct VCC	A temporary bidirectional point-to-point VCC set up by the LEC to the LECS.
Maximum Frame Size	The length (in bytes) of the largest frame field. Can be **1516**, **1580**, **4544**, **9234**, or **18190**.
Control Direct VCC	A bidirectional point-to-point VCC set up by the LEC to the LES.
Control Distribute VCC	A unidirectional point-to-multipoint VCC set up by the LES to the LECs in the ELAN.
Multicast Send VCC	A bidirectional point-to-point VCC set up by the LEC to the BUS for sending multicast data to the BUS.
Multicast Forward VCC	A unidirectional point-to-multipoint VCC set up from the BUS to the LECs in the ELAN.
Last Error	The last error experienced by this LEC, if applicable. This field is reset when the switch is rebooted.

[1] When you change the administrative status from `down` to `up`, it takes a few seconds for the change to occur and to be reflected in the `OperStatus` field. Therefore, it is possible for the `AdminStatus` to be `up`, but the `OperStatus` to be `down`. If you refresh the display, the change has taken place and is reflected here.

LANE Configuration Commands

13.4   LECS Configuration Commands

These commands allow you to configure the LAN Emulation Configuration Server (LECS). You can display the list of available subcommands by typing **lecs ?** at the **lane** level.

```
myswitch::configuration lane> lecs ?
  admin           delete          new             show
  get
```

13.4.1 Configuring the LECS Administrative Status

This command lets you change the administrative status of the LECS to up (start a LECS service) or down (stop a LECS service). Enter the following parameters:

```
myswitch::configuration lane lecs> admin <LECS index | LECS index range (x-y)>
(up | down)
```

13.4.2 Deleting a LECS

This command lets you delete (stop) a specified LECS service. Enter the following parameters:

```
myswitch::configuration lane lecs> delete <LECS index | LECS index range (x-y)>
```

The parameters for admin and delete are defined as follows:

Parameter	Description
LECS index	The unique integer index of the LECS that is dynamically assigned by AMI when a LECS is created to identify this service from any other service in the same class. Found under the `Index` field using the **conf lane lecs show** command.
LECS index range (x-y)	The range of index numbers of the LECSs that you want to start or stop. The index is found under the `Index` field using **conf lane lecs show**.
up|down	**up** changes the administrative status of the designated LECS index to up. **down** changes the administrative status of the designated LECS index to down.

13.4.3 Creating a LECS

This command lets you create (start) a LECS service. Enter the following parameters:

```
myswitch::configuration lane lecs> new <LECS Selector byte (HEX)>
[-db <LECS database file>]
[-default <LES atm address>] [<LECS-wka> | none]
```

The parameters for new are defined as follows:

Parameter	Description
LECS Selector byte (HEX)	The 20th byte of the ATM address of the host or switch that is to run a LECS service (entered in hexadecimal format). Use the `conf atmarp getnsap` command to display the entire ATM address.
-db <LECS database file>	The full path to the location and name of the LECS database file. The default file is `lecs.cfg`. For information about configuring this file, refer to the *ATM Switch Network Configuration Manual.*
-default <LES atm address>	The default LES address to use in case the LECS configuration file is inaccessible.
<LECS-wka> | none	`none` means the well-known address is disabled so that the LECS can only be contacted by using the switch's actual address (with selector byte). `LECS-wka` means that the ATM address that you enter is going to be used as the well-known address instead of the ATM Forum well-known address.

NOTE ▶ Although more than one LECS can be created on a switch, you must ensure that no two LECS are listening on the same address.

NOTE ▶ If you create a LECS and then change the NSAP prefix of the switch at a later time, the change is NOT automatically reflected in the LECS's address. You must delete the LECS and recreate it using the new NSAP prefix.

LANE Configuration Commands

13.4.4 Displaying LECS Information

This command lets you display the current LECS information. To display information about the current LECS configuration on the switch, enter the following parameters:

```
myswitch::configuration lane lecs> show
  Index  AdminStatus  OperStatus   Selector  WKA         Database
      1  up           up               0x01  atm-forum   lecs.cfg
         Default LES :0x47.0005.80.ffe100.0000.f21a.096b.0020481a096b.a2
      2  up           up               0x02  none        lecs.cfg
      3  up           up               0x03  other       lecs.cfg
         wka: 0x47.0005.80.ffe100.0000.f21a.11f1.0020481a11f8.99
```

The fields in this display are defined as follows:

Field	Description
Index	The unique index number of the LECS that identifies it from other LECSs. This number is dynamically assigned by the switch when the LECS is created.
AdminStatus	Reflects any changes that you have made to the status of the LECS. **up** means you have started the LECS. **down** means you have stopped the LECS.
OperStatus[1]	Reflects the actual current status of the LECS. **up** means the LECS is currently active. **down** means the LECS is currently inactive.
Selector	The selector byte portion (20th byte) of the ATM address, in hexadecimal format, of the host or switch that is running the LECS.
WKA	The state of the well-known address for this particular LECS. **atm-forum** means the LECS is using the default address as defined by the ATM Forum. **none** means the well-known address has been disabled. **other** means the well-known address has been redefined, in which case the new address appears on the following line.
Database	The full path to the location and name of the LECS database file. The default file is `lecs.cfg`.
Default LES	The default LES address to use in case the LECS configuration file is inaccessible. If a default LES address has not been specified, this field is not displayed.

[1.] When you change the administrative status from `down` to `up`, it takes a few seconds for the change to occur and to be reflected in the `OperStatus` field. Therefore, it is possible for the `AdminStatus` to be `up`, but the `OperStatus` to be `down`. If you refresh the display, the change has taken place and is reflected here.

To display information about a particular LECS that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane lecs> show [<LECS index>]
```

For example, to display information about the LECS with an index number of 1, enter the following parameters:

```
myswitch::configuration lane lecs> show 1
 Index   AdminStatus  OperStatus   Selector  WKA        Database
     1   up           up             0x01    atm-forum  lecs.cfg
         Default LES :0x47.0005.80.ffe100.0000.f21a.096b.0020481a096b.a2
```

The fields in this display are defined in the same manner as those listed above in the example for all of the LECS configured on the switch.

> **NOTE**
>
> If there is an error in the LECS.CFG file that prevents the LECS from starting, an error is printed to the console (if the console is enabled) indicating the line number in which the error occurred.

13.4.5 Getting the LECS Configuration File

This command lets you download the LECS configuration file. Enter the following:

```
myswitch::configuration lane lecs> get <host>:<remotefile> [<localfile>]
```

These parameters are defined as follows:

Parameter	Description
host	The IP address of the host from which the LECS database file is to be retrieved.
remotefile	The name of the LECS database file that is to be retrieved.
localfile	The name of the file where the retrieved LECS database file is to be stored. The default local file is `lecs.cfg`.

NOTE For information about configuring this file, refer to the *ATM Switch Network Configuration Manual.*

If you have configured the transfer protocol to be FTP using **conf system protocol**, you only need to enter the command shown above to retrieve the LECS file. After you enter the command shown above, you are prompted for the remote userid and password of the remote host from which you are retrieving the file. For information about the **conf system protocol** command, see Part 2 of the *AMI Configuration Configuration Commands Reference Manual.*

If you have configured the transfer protocol to be TFTP (this is the default) using **conf system protocol**, the remote host from which the LECS file will be retrieved must be running the TFTP server code. If you are unsure of how to do this, see the *ATM Switch Installation and Maintenance Manual.*

13.5   LES Configuration Commands

These commands allow you to configure the LAN Emulation Server (LES). You can display
the list of available subcommands by typing **les ?** at the **lane** level.

```
myswitch::configuration lane> les ?
  admin            delete           new                  peeradd
  peerdelete       security         show
```

13.5.1 Configuring the LES Administrative Status

This command lets you change the administrative status of the LES to up (start a LES service)
or down (stop a LES service). Enter the following parameters:

```
myswitch::configuration lane les> admin <LES index | LES index range (x-y)> (up | down)
```

13.5.2 Deleting a LES

This command lets you delete a specified LES. Enter the following parameters:

```
myswitch::configuration lane les> delete <LES index | LES index range (x-y)>
```

The parameters for admin and delete are defined as follows:

Parameter	Description
LES index	The unique integer index of the LES that is dynamically assigned by AMI when a LES is created to identify this service from any other service in the same class. Found under the Index field using the **conf lane les show** command.
LES index range (x-y)	The range of index numbers of the LES that you want to start or stop. The index is found under the Index field using **conf lane les show**.
up | down	**up** changes the administrative status of the designated LES index to up. **down** changes the administrative status of the designated LES index to down.

13.5.3 Creating a LES

This command lets you create a new LES. Enter the following parameters:

```
myswitch::configuration lane les> new <LES Selector Byte (HEX)> <LES name> \
        [-bus <BUS Selector Byte (HEX)>] \
        [-type (ethernet | token-ring)] \
        [-mtu (1516 | 1580 | 4544 | 9234 | 18190)] \
        [-secure wka | <LECS ATM Address>] \
        [-registertlvs (enable | disable)] \
        [-elanid <ELAN ID>] \
        [-segid <TR Segment ID>] \
        [-fwdarp (enable | disable)] \
        [-anycast <LES Anycast ATM Address>] \
        [-peers <atm-addr> ...]
```

The parameters for new are defined as follows:

Parameter	Description
LES Selector Byte (HEX)	The 20th byte of the ATM address of the switch that is to run a LES service (entered in hexadecimal format). This command starts a co-located BUS, which means that the LES and BUS services for a particular ELAN are running on the same switch. You cannot create a BUS separately. By default, the LES and BUS will use the same selector byte.
BUS Selector Byte (HEX)	The 20th byte of the ATM address of the BUS associated with this LES. Use this option only if you need to have the BUS use a different selector byte than the LES.
LES name	The user-defined name that helps you to remember what ELAN this LES services.
-type (ethernet | token-ring)	`ethernet` means the LAN type is Ethernet. `token-ring` means the LAN type is Token Ring. The default is `ethernet`.
-mtu (1516 | 1580 | 4544 | 9234 | 18190)	The length (in number of bytes) of the largest frame field (or MTU size). The default for Ethernet is `1516`. The default for Token Ring is `4544`. The MTU size must match that of the other hosts and edge devices on the ELAN.
-secure wka	<LECS ATM Address>	Indicates you want to activate ELAN access control. If you enter `wka`, the ATM Forum well-known address is used when running with ELAN access control enabled. In this case, you do not have to type the actual well-known address. If you are using an LECS address that is different than the well-known address, then you must type the full LECS ATM address to be used when running with ELAN access control enabled. ELAN access control is disabled by default.
-registertlvs (enable | disable)	`enable` means the LES supports MPOA operation by accepting flow descriptors from LANE/MPOA clients (LEC/MPCs) that register with it and distributing these flow descriptors to LEC/MPCs in response to LE-ARP queries. `disable` means the LES collects flow descriptors parameters from LEC/MPCs that register with it, but does not distribute these flow descriptors to LEC/MPCs in response to LE-ARP queries. The default is `enable`.

Parameter	Description
-elanid <ELAN ID>	The unique ID for this ELAN. The LES gives this value to the LANE clients in its join response and the clients use it in their data frames. Allowable values are any positive integers.
-segid <TR Segment ID>	The segment ID for a Token Ring ELAN that is used by Token Ring ATM bridges in source routing. The LES gives this information to the LANE clients in its join response. Allowable values are from 0 - 4096.
-fwdarp (enable | disable)	`enable` means the LES forwards all of the LE_ARP requests to all of the LANE clients, including non-proxy clients. The LES never answers any LE_ARP requests by itself. `disable` means the LES answers LE_ARP requests for those LANE clients that registered an address with it. For all unregistered addresses, the LE_ARP requests are only forwarded to proxy clients. The default is `disable`.
-anycast <LES Anycast ATM Address>	The unique anycast ATM address that all LECs in a given ELAN use to connect to one of the Distributed LAN Emulation (DLE) servers for that ELAN. If this option is not used, then the LES address is used, meaning that this is a stand-alone LES/BUS pair (i.e., DLE services are not running). The anycast address <u>must</u> be unique within the first 19 bytes. Anycast addresses cannot be all zeros and the first 19 bytes cannot be all zeros. Anycast addresses cannot be the same as any ATM Forum-defined anycast address.
-peers <atm-addr>	The unique ATM address(es) of this DLE server's peer(s). It is through these addresses that all servers in a DLE ELAN connect to and communicate with each other. You must explicitly enter the local LES ATM address in the list of DLE peers. It does not matter in which order the local address appears in the list.

NOTE

If you create a LES and then change the NSAP prefix of the switch at a later time, the change is NOT automatically reflected in the LES's address. You must delete the LES and recreate it using the new NSAP prefix. If you are using DLE, you must also change the LES's address in the lists of all of the other switches that are running DLE peers. You can use the `conf lane les peeradd` (described in Section 13.5.4) and the `conf lane les peerdelete` command (described in Section 13.5.5) to do this.

LANE Configuration Commands

13.5.4 Adding a DLE Peer

This command lets you add a DLE peer server to the list of DLE peer servers for an existing DLE ELAN without destroying and recreating the existing peers. Enter the following:

```
myswitch::configuration lane les> peeradd <LES index> <Peer ATM Address>
```

These parameters are defined as follows:

Parameter	Description
LES index	The index number of the LES that you want to add to the list of DLE peer servers for a particular ELAN. This index number is dynamically assigned by the switch when the LES is created. It can be found under the Index field using the **conf lane les show** command.
Peer <ATM Address>	The unique ATM address of the DLE peer server you want to add. It is through this address that all of the DLE peer servers in a DLE ELAN connect to and communicate with each other.

> **NOTE** When using this command, the LES must already exist and the ELAN that it services must be a DLE ELAN.

The LES is temporarily administered down while it is being added to the list of peers and comes back up once the process is complete. Because of this, a warning is shown when you issue this command and you are asked if you want to continue. For example:

```
myswitch::configuration lane les> peeradd 1
0x47.0005.80.ffe100.0000.f21a.23c0.0020481a23c0.75

WARNING: This command will disable the LES service for the
duration of execution. Full ELAN connectivity may not be restored
until all the LANE clients reconnect to the LES service

Do you want to continue [n]? y
```

Type **n** or **<ENTER>** to abort the command or type **y** to continue adding the peer.

13.5.5 Deleting a DLE Peer

This command lets you delete a DLE peer server from the list of DLE peer servers for an existing ELAN without destroying and recreating the existing peers. Enter the following:

```
myswitch::configuration lane les> peerdelete <LES index> <Peer ATM Address>
```

These parameters are defined as follows:

Parameter	Description
LES index	The index number of the LES that you want to delete from the list of DLE peer servers for a particular ELAN. This index number is dynamically assigned by the switch when the LES is created. It can be found under the `Index` field using the **conf lane les show** command.
Peer <ATM Address>	The unique ATM address of the DLE peer server you want to delete. It is through this address that all of the DLE peer servers in a DLE ELAN connect to and communicate with each other.

NOTE ▶ When using this command, the LES must already exist and the ELAN that it services must be a DLE ELAN.

The LES is temporarily administered down while it is being deleted from the list of peers and comes back up once the process is complete. Because of this, a warning is shown when you issue this command and you are asked if you want to continue. For example:

```
myswitch::configuration lane les> peerdelete 2
0x47.0005.80.ffe100.0000.f21a.23c0.0020481a23c0.65

WARNING: This command will disable the LES service for the
duration of execution. Full ELAN connectivity may not be restored
until all the LANE clients reconnect to the LES service

Do you want to continue deleting the peer [n]? y
```

Type **n** or **<ENTER>** to abort the command or type **y** to continue deleting the peer.

NOTE ▶ The local DLE peer server address cannot be deleted using this command.

13.5.6 Enabling/Disabling ELAN Access Control

This command lets you enable or disable ELAN access control for an existing ELAN without
destroying and recreating the LES. (For more information about ELAN access control, see the
ATM Switch Network Configuration Manual.) Enter the following:

```
myswitch::configuration lane les> security <LES index> (disable|enable)
[<wka|LECS Addr>]
LECS address required for enabling security
```

These parameters are defined as follows:

Parameter	Description
disable | enable	**disable** means that you want to de-activate ELAN access control on this LES. **enable** means that you want to activate ELAN access control on this LES.
wka | LECS Addr	You must enter one of these parameters if you want to enable ELAN access control. If you enter **wka**, the ATM Forum well-known address is used when running with ELAN access control enabled. In this case, you do not have to type in the actual well-known address, just enter **wka**. If you are using an LECS address that is different than the well-known address, then you must type the full LECS ATM address to be used when running with ELAN access control enabled.

NOTE ➤ When using this command, the LES must
already exist.

The LES is temporarily administered down while ELAN access control is being enabled or dis-
abled and comes back up once the process is complete. Because of this, a warning is shown
when you issue this command and you are asked if you want to continue. For example:

```
myswitch::configuration lane les> security 2 enable
0x47.0005.80.ffe100.00005.f21a.23c0.00200481a23c0.00

WARNING: This command will disable the LES service for the
duration of execution. Full ELAN connectivity may not be restored
until all the LANE clients reconnect to the LES service

Do you want to continue deleting the peer [n]? y
```

Type **n** or **<ENTER>** to abort the command or type **y** to continue the process.

13.5.7 Displaying LES Information

This command lets you display the current LES information. To display information about every LES that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane les> show
      Admin Oper  Les                      ELAN                                  ELAN
 Index State State Sel  Type      MTU   Name    SECURE   FwdTLVs FwdARPs ID
 1     up    up    0x10 ethernet  1516  elan1   disable enable  disable 100
        LES : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
        BUS : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10 (Co-Located)
            : c5000580ffe1000000f21c10bb0020323c10bb10 (ANYCAST)
        LECS  :0x47.0079.00.000000.0000.0000.0000.00a03e000001.00
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
 2     up    up    0x20 ethernet  1516  elan1   disable enable  enable  200
        LES : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        BUS : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10 (Co-Located)
            : c5000580ffe1000000f21c10bb0020323c10bb10 (ANYCAST)
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
 3     up    up    0x30 ethernet  1516  elan1   disable enable  enable  0
        LES : 0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
        BUS : 0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70 (Co-Located)
            : c5000580ffe1000000f21c10bb0020323c10bb10 (ANYCAST)
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        PEER  :0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
 4     up    up    0x40 token-ring 4544 elan2   disable enable  disable 300
        Token-Ring segment ID : 1000
        LES : 0x47.0005.80.ffe100.0000.f21a.24aa.0020481a6754.33
        BUS : 0x47.0005.80.ffe100.0000.f21a.24aa.0020481a6754.33 (Co-Located)
 5     up    up    0x50 token-ring 4544 elan2   disable enable  enable  0
        Token-Ring segment ID : 2000
        LES : 0x47.0005.80.ffe100.0000.f21c.26bb.0020481b4852.44
        BUS : 0x47.0005.80.ffe100.0000.f21c.26bb.0020481b4852.44 (Co-Located)
Press return for more, q to quit: q
```

The fields in this display are defined as follows:

Field	Description
Index	The unique index number of the LES that identifies it from other LESs. This number is dynamically assigned by the switch when the LES is created.
Admin State	Reflects any changes that you have made to the status of the LES. **up** means you have started the LES. **down** means you have stopped the LES.
Oper State	Reflects the actual current status of the LES. **up** means the LES is currently active. **down** means the LES is currently inactive. If you change the administrative status from down to up, it takes a few seconds for the change to occur and to be reflected in the OperStatus field. Therefore, it is possible for the AdminStatus to be up, but the OperStatus to be down. If you refresh the display, the change has taken place and is reflected here.
Les Sel	The selector byte portion (20th byte) of the ATM address, in hexadecimal format, of the host or switch that is the LES. This address is used by the peers of this DLE server when they need to contact it.
Type	Shows what type of ELAN this is. **ethernet** means that the LAN type is Ethernet. **token-ring** means that the LAN type is Token Ring.
MTU	The maximum data frame size. Can be **1516**, **4544**, **9234**, or **18190**. The default for Ethernet is **1516**. The default for Token Ring is **4544**.
ELAN Name	The name of the ELAN that this LES serves.
SECURE	**enable** means that ELAN access control is running. **disable** means that ELAN access control is not running.
FwdTLVs	**enable** means the LES accepts TLV parameters from LEC/MPCs that register with it and distributes these TLVs to LEC/MPCs in response to LE-ARP queries. **disable** means the LES collects TLV parameters from LEC/MPCs that register with it, but does not distribute these TLVs to LEC/MPCs in response to LE-ARP queries.
FwdARPs	**enable** means the LES forwards all of the LE_ARP requests to all of the LANE clients, including non-proxy clients. The LES never answers any LE_ARP requests by itself. **disable** means the LES answers LE_ARP requests for those LANE clients that registered an address with it. For all unregistered addresses, the LE_ARP requests are only forwarded to proxy clients.
ELAN ID	The unique ID for this ELAN. The LES gives this value to the LANE clients in its join response and the clients use it in their data frames. 0 is displayed if the user did not configure a value.
Token-Ring segment ID	The segment ID for a Token Ring ELAN that is used by Token Ring ATM bridges in source routing. The LES gives this information to the LANE clients in its join response.
LES	The full ATM address of the LES, which the DLE peers use to connect to this server. (If DLE is not running, the LECs that want to join the ELAN use this address to contact the LES.)
BUS	The full ATM address of the BUS for the ELAN.
(ANYCAST)	The unique anycast ATM address that all LECs in a given ELAN use to connect to one of the DLE servers for that ELAN. If no address is provided by the user, then this is a stand-alone LES/BUS pair (i.e., DLE services are not running in that ELAN).

Field	Description
LECS	The LECS address. This address is displayed only if ELAN access control has been enabled.
PEER	If there are any DLE peer servers, shows the full ATM address of each.

To display information about a particular LES that is currently configured on the switch, enter the following parameters:

```
myswitch::configuration lane les> show [<LES index>] [(advanced)]
```

For example, to show information about the LES with an index number of 2, enter the following:

```
       Admin Oper  Les                      ELAN                             ELAN
 Index State State Sel  Type       MTU  Name       SECURE  FwdTLVs FwdARPs ID
 2     up    up    0x20 ethernet   1516 elan1      disable enable  enable  200
        LES : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        BUS : 0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10 (Co-Located)
            : c5000580ffe1000000f21c10bb0020323c10bb10 (ANYCAST)
        PEER   :0x47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
        PEER   :0x47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
        PEER   :0x47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
```

The fields in this display are defined in the same way as those listed in the previous example.

Enter the following parameters to display advanced information about every LES that is currently configured on the switch. Information similar to the following is displayed when DLE is running:

```
myswitch::configuration lane les> show advanced
ELAN Name: "elan1"
  LES:    47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
          c5000580ffe1000000f21c10bb0020323c10bb10 (anycast)
          Point-to-Multipoint VCC to Peers: 0.714
  BUS:    47.0005.80.ffe100.0000.f21c.10bb.0020481c10bb.10
  Peer #1: 47.0005.80.ffe100.0000.f21c.10bb.0020481a42bb.10
          Point-to-Point VCC to Peer: 0.713
          Point-to-Multipoint VCC from Peer: 0.61
  Peer #2: 47.0005.80.ffe100.0000.f21c.10bb.0020482a55a0.70
          Point-to-Point VCC to Peer: 0.711
          Point-to-Multipoint VCC from Peer: 0.62
  LAN Type: Ethernet/IEEE 802.3    Maximum Data Frame Size: 1516
  Non-proxy Control Distribute VCC: 0.716
      Proxy Control Distribute VCC: -.-
            Multicast Forward VCC: 0.718
  Number of local clients: 4
LEC #1 at 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.00 (non-proxy)
  00-20-48-10-13-f2 -> 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.00
  Control Direct VCC: 0.715
LEC #2 at 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.01 (non-proxy)
  02-20-48-10-13-f2 -> 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.01
  Control Direct VCC: 0.719
LEC #3 at 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.02 (non-proxy)
  06-20-48-10-13-f2 -> 47.0005.80.ffe100.0000.f21a.341a.0020481013f2.02
  Control Direct VCC: 0.721
LEC #4 at 47.0005.80.ffe100.0000.f21a.341a.0020481a341a.00 (non-proxy)
  00-20-48-1a-34-1a -> 47.0005.80.ffe100.0000.f21a.341a.0020481a341a.00
  Control Direct VCC: 0.724
```

The fields in this display are defined as follows:

Field	Description
ELAN Name	The names of any ELANs associated with this switch.
LES	When DLE is configured, the first ATM address displayed is the LES unicast address that is used by the other DLE peer servers to connect to this server. The second is the LES any-cast address that is used by LECs to connect to the ELAN. (All DLE peer servers in an ELAN must use the same anycast address.) The Point-to-Multipoint VCC to Peers is the VCC on which packets are transmitted by this LES when they are flooded to all of the DLE peer servers. When DLE is not configured, only the single ATM address of the LES that services this particular ELAN is displayed.
BUS	The ATM address of the BUS that services the LECs that are directly connected to this DLE peer server.
Peer #1	The address of one of the DLE peer servers. This is the address at which the LES tries to connect to this DLE peer server. The Point-to-Point VCC to Peer is the VCC that is used to send or receive packets to or from this particular DLE peer server. The Point-to-Multipoint VCC from Peer is the VCC on which packets are received by the LES from this DLE peer server when they are flooded to all of the DLE peer servers.
LAN Type	The type of emulated LAN. Either `Ethernet/IEEE 802.3` or `TokenRing/IEEE 802.5`.
Maximum Data Frame Size	The length (in bytes) of the largest frame field. Can be **1516**, **1580**, **4544**, **9234**, or **18190**.
Non-proxy Control Distribute VCC	The point-to-multipoint connection that the LES maintains to all of the non-proxy LECs that it services.
Proxy Control Distribute VCC	The point-to-multipoint connection that the LES maintains to all of the proxy LECs that it services. This entry is blank if no proxy LECs have joined this ELAN.
Multicast Forward VCC	The VCC on which the BUS forwards broadcast and multicast traffic to all the LECs.
Number of local clients	This is the number of clients connected to this server.
LEC	The LEC ATM address, shows the MAC-to-ATM address mapping for each LEC in this ELAN, and shows whether or not this LEC is a proxy.
Control Direct VCC	The point-to-point connection that the LES maintains to this particular LEC.

When DLE is not running, the same information as above is displayed, except there is no any-cast nor peer information to be displayed.

CHAPTER 14 **Network Module Commands**

These commands let you configure the network modules in your switch fabric. The list of available subcommands is displayed by typing **module ?** at the **configuration** level.

```
myswitch::configuration> module ?
  admin           fram>           reset           show
  traffic>
```

NOTE The **fram>** commands are not available on an ASX-200WG, LE 25, LE 155, nor an ASX-4000.

14.1   Configuring the State of a Network Module

This command allows you to change the state of a network module up or down. This command should only be used when you want to test the network module. Enter the following parameters:

```
myswitch::configuration module> admin <module> (up | down)
```

These parameters are defined as follows:

Parameter	Description
module	The network module that is to be started or stopped.
up | down	**up** brings that network module back on-line. **down** causes the designated network module to be be taken off-line temporarily so that it can be tested.

NOTE

This command has no effect on Circuit Emulation Services (CES) DS1 and E1 network modules. You can only administer the individual ports on a CES module up or down using the **conf port cesds1 admin** or **conf port cese1 admin** command.

When the **conf module admin** *<module>* **down** command is entered, a warning message is displayed as follows:

```
myswitch::configuration module> admin 2B down

Disabling the network module will destroy all existing connections
going through it.

Disable the network module [n]?
```

Entering **n** or pressing **<ENTER>** aborts the command. Entering **y** tears down all of the existing connections and temporarily places the network module out of service.

To place the module back into service, use the command **conf module admin** *<module>* **up**. At that point, PVCs that are stored in the configuration database are re-established and SVCs are dynamically re-established.

14.2   Configuring a *FramePlus* Network Module

These commands allow you to configure a *FramePlus* network module. The list of available subcommands is displayed by typing **?** at the **fram** level.

```
myswitch::configuration module fram> ?
  application       setmem            highthresholds    lowthresholds
  oamf5rate         upgrade           oamf5             stats
  show
```

14.2.1 Configuring an Application

This command lets you configure a *FramePlus* network module to run either a Frame Relay over ATM or a frame-based User to Network Interface (FUNI) application. The network module can only run one application or the other at a time. Enter the following:

```
myswitch::configuration module fram> application <module> <application_key>
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
application_key	Indicates the key associated with the type of service to be used on this *FramePlus* network module. *FramePlus* network modules run Frame Relay by default. To run FUNI services, please contact FORE's Technical Assistance Center for a FUNI key. To change from FUNI back to Frame Relay, the key is **fratm170358**. You can look at the Appln field under **conf mod show** to see which application is currently running.

The application key should be specified only if you want to reconfigure the network module to run a different type of application. Before changing the application, you must administer the network module down as follows:

```
myswitch::configuration module> admin 4a down
Disabling the network module will destroy all
existing connections going through it.
Disable the network module [n]? y
```

Then, you can change the application. When you change the application, the switch deletes all existing services and PVCs that use a different application, and removes them from the CDB (i.e., if you are changing from Frame Relay to FUNI, the switch deletes existing Frame Relay information, and vice versa). The switch warns you as follows:

```
myswitch::configuration module fram> app 4a XXXXXXXXXX
Changing application may cause deletion of service/connection on switch and CDB.
Proceed [n]? y
```

Then, you need to administer the network module up as follows:

```
myswitch::configuration module> admin 4a up
```

14.2.2 Configuring the Buffer on a *FramePlus* Network Module

This command lets you configure the high and low priority EPD/PPD buffer sizes on a *FramePlus* network module. The total buffer size is 32,768 cells. There are four fixed configuration models for partitioning the buffer between the two priorities, as shown in the table below. You can select one of these.

> **NOTE** See Appendix D in the *ATM Switch Network Configuration Manual* for more information about configuring the buffer.

Enter the following:

```
myswitch::configuration module fram> setmem <module> (highzero | high1quarter |
high2quarter | high3quarter)
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
highzero	The high priority buffer will have 0 cells and the low priority buffer will have 32,768 cells. (This implies that the **highthresholds** command under **conf module fram** will have no meaning if it is used.) This is the default value.
high1quarter	The high priority buffer will have 8,192 cells and the low priority buffer will have 24,576 cells.
high2quarter	The high priority buffer will have 16,384 cells and the low priority buffer will have 16,384 cells.
high3quarter	The high priority buffer will have 24,576 cells and the low priority buffer will have 8,192 cells.

Before you can change the buffer allocation, you must first take the network module out of service by administering it down as follows:

```
myswitch::configuration module> admin 1d down
Disabling the network module will destroy all
existing connections going through it.

Disable the network module [n]? y

myswitch::configuration module> fram
myswitch::configuration module fram> set 1d high3quarter
```

14.2.3 Configuring the High Priority Buffer

This command lets you configure EPD/PPD thresholds for the high priority buffer on a *FramePlus* network module.

> **NOTE** ▶ This command has no effect when the `highzero` option has been selected under `conf module fram setmem`.

> **NOTE** ▶ See Appendix D in the *ATM Switch Network Configuration Manual* for more information about configuring this option.

Enter the following:

```
myswitch::configuration module fram> highthresholds <module>
          [-clp0epd (50|62|75|87)]
          [-clp1ppd (50|62|75|87)]
          [-clp1epd (25|37|50|62)]
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
clp0epd	The threshold for CLP0EPD, in percentage of the available size, for the high priority queue buffer. The default value is 62%. EPD discards packets once this CLP=0 threshold has been surpassed.
clp1ppd	The threshold for CLP1PPD, in percentage of the available size, for the high priority queue buffer. The default value is 50%. PPD discards cells once this CLP=1 threshold has been surpassed.
clp1epd	The threshold for CLP1EPD, in percentage of the available size, for the high priority queue buffer. The default value is 37%. EPD discards packets once this CLP=1 threshold has been surpassed.

> **NOTE** ▶ There is a CLP0PPD threshold that is automatically set for you as 87.5% of the buffer size and cannot be changed. (It is not displayed in AMI, but you need to know this value because the other thresholds are calculated on the <u>remaining</u> buffer size.)

14.2.4 Configuring the Low Priority Buffer

This command lets you configure EPD/PPD thresholds for the low priority buffer on a *FramePlus* network module.

NOTE   See Appendix D in the *ATM Switch Network Configuration Manual* for more information about configuring this option.

Enter the following:

```
myswitch::configuration module fram> lowthresholds <module>
          [-clp0epd (50|62|75|87)]
          [-clp1ppd (50|62|75|87)]
          [-clp1epd (25|37|50|62)]
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
clp0epd	The threshold for CLP0EPD, in percentage of the available size, for the low priority queue buffer. The default value is 62%. EPD discards packets once this CLP=0 threshold has been surpassed.
clp1ppd	The threshold for CLP1PPD, in percentage of the available size, for the low priority queue buffer. The default value is 50%. PPD discards cells once this CLP=1 threshold has been surpassed.
clp1epd	The threshold for CLP1EPD, in percentage of the available size, for the low priority queue buffer. The default value is 37%. EPD discards packets once this CLP=1 threshold has been surpassed.

NOTE   There is a CLP0PPD threshold that is automatically set for you as 87.5% of the buffer size and cannot be changed. (It is not displayed in AMI, but you need to know this value because the other thresholds are calculated on the <u>remaining</u> buffer size.)

14.2.5 Configuring OAM Rates

This command lets you set the receive and transmit rates of AIS and RDI cells on a *FramePlus* network module. Enter the following:

```
myswitch::configuration module fram> oamf5rate <module> [-aisrx <sec>]
[-aistx <sec>] [-rdirx <sec>] [-rditx <sec>]
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
aisrx <sec>	Time interval, in seconds, at which an AIS OAM cell is to be received. The default is **3**.
aistx <sec>	Time interval, in seconds, at which an AIS OAM cell is to be transmitted. The default is **1**.
rdirx <sec>	Time interval, in seconds, at which an RDI OAM cell is to be received. The default is **3**.
rditx <sec>	Time interval, in seconds, at which an RDI OAM cell is to be transmitted. The default is **1**.

14.2.6 Upgrading the Software on a *FramePlus* Network Module

This command lets you upgrade the software on a *FramePlus* network module. For more information, see Appendix D in the *ATM Switch Network Configuration Manual*. Enter the following:

```
myswitch::configuration module fram> upgrade <module> <remotehost>:<fullpath to
remotefile>
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module on which you want to upgrade the software.
remotehost	The IP address of the remote host on which the upgrade file resides.
full path to remotefile	The full path name of the upgrade file.

Since this command uses TFTP as the transfer protocol, the remote host on which the upgrade file resides must be a tftpboot server. If you are unsure of how to configure the bootp server and the tftpboot server properly, see Chapter 4 of the *ATM Switch Installation and Maintenance Manual* for your switch.

14.2.7 Configuring OAM F5 Functionality

This command lets you enable or disable the generation of OAM F5 cells at the module level on a *FramePlus* network module.

NOTE ▶ This command only affects Frame Relay service on a *FramePlus* network module. It does not have any effect when the FUNI application is being used.

Enter the following:

```
myswitch::configuration module fram> oamf5 <module> (enabled|disabled)
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
enabled | disabled	**enabled** indicates that OAM F5 (path) cells will be generated on the specified *FramePlus* network module. **disabled** indicates that OAM F5 (path) cells will not be generated on the specified *FramePlus* network module. The default is **enabled**.

14.2.8 Enabling and Disabling Statistics

This command lets you enable or disable the collection of statistics on a *FramePlus* network module. Enter the following:

NOTE ▶ To collect Frame Relay statistics, they must be enabled both here <u>and</u> at the Frame Relay service level under **conf fratm stats**. To collect FUNI statistics, they must be enabled both here <u>and</u> at the FUNI service level under **conf funi stats**. They are enabled here by default.

```
myswitch::configuration module fram> stats <module> (enabled|disabled)
```

These parameters are defined as follows:

Parameter	Description
module	The *FramePlus* network module that is to be configured.
enabled | disabled	**enabled** means statistics are collected on the designated *FramePlus* network module. **disabled** means statistics are not collected on the designated *FramePlus* network module. The default is **enabled**. When statistics are disabled, the counters show zeros.

14.2.9 Displaying the Settings on a *FramePlus* Network Module

This command lets you display information about the configuration of a *FramePlus* network module. Enter the following:

```
myswitch::configuration module fram> show
Module Appln Appl   Boot   Stats    Oam     Operational      Product
             swRel  swRel  monitor  monitor state            number
1A      fratm 1.2.0 1.0.0  enabled  enabled appluprunning    NMFR-4/E1A
1C      fratm 1.2.0 1.0.0  enabled  enabled appluprunning    NMFR-4/E1A
```

These fields are defined as follows:

Field	Description
Module	The *FramePlus* network module that has been configured.
Appln	The current application that is running on this *FramePlus* network module. Displays `fratm` for Frame Relay or `funi` for FUNI. The default application is `fratm`.
Appln swRel	The embedded software release that is currently running on this *FramePlus* network module.
Boot swRel	The boot software release on this *FramePlus* network module.
Stats monitor	Shows if statistics collection is enabled (on) or disabled (off) for this *FramePlus* network module.
Oam monitor	Shows if OAM F5 cell generation is enabled (on) or disabled (off) for this *FramePlus* network module.
Operational state	Shows the operational status of the current application.
Product number	The FORE Systems product number for this module.

You can display configuration information for an individual *FramePlus* network module as follows:

```
myswitch::configuration module fram> show 1a
Module Appln Appl   Boot   Stats    Oam     Operational      Product
             swRel  swRel  monitor  monitor state            number
1A      fratm 1.2.0 1.0.0  enabled  enabled appluprunning    NMFR-4/E1A
```

The fields in this display are defined in the same manner as those in the previous example.

If no *FramePlus* network modules are installed, then the following is displayed:

```
No FRAM netmod configuration information is available
```

You can also display advanced information about the configuration of a *FramePlus* network module. You cannot enter only the advanced option. You must enter the specific module <u>and</u> the advanced option as follows:

```
myswitch::configuration module fram> show [<module>] [advanced]
myswitch::configuration module fram> show 1a advanced
Module Buffer Buffer Clp0Epd  Clp1Ppd  Clp1Epd   AIS  AIS  RDI  RDI
        high   low   high low high low high low   rx   tx   rx   tx
1A         0 32768 N/A   62 N/A   50 N/A   37     3    1    3    1
```

These fields are defined as follows:

Field	Description
Module	The *FramePlus* network module that has been configured.
Buffer high	The used-configured size for the high priority buffer, in cells. Can be one of four sizes.
Buffer low	The used-configured size for the low priority buffer, in cells. Can be one of four sizes.
Clp0Epd high	The CLP0EPD threshold for the high priority buffer, in percentage of the available size. The default value is 62%. EPD discards packets once this CLP=0 threshold has been surpassed.
Clp0Epd low	The CLP0EPD threshold for the low priority buffer, in percentage of the available size. The default value is 62%. EPD discards packets once this CLP=0 threshold has been surpassed.
Clp1Ppd high	The CLP1PPD threshold for the high priority buffer, in percentage of the available size. The default value is 50%. PPD discards cells once this CLP=1 threshold has been surpassed.
Clp1Ppd low	The CLP1PPD threshold for the low priority buffer, in percentage of the available size. The default value is 50%. PPD discards cells once this CLP=1 threshold has been surpassed.
Clp1Epd high	The CLP1EPD threshold for the high priority buffer, in percentage of the available size. The default value is 37%. EPD discards packets once this CLP=1 threshold has been surpassed.
Clp1Epd low	The CLP1EPD threshold for the low priority buffer, in percentage of the available size. The default value is 37%. EPD discards packets once this CLP=1 threshold has been surpassed.
AIS rx	Time interval, in seconds, at which an AIS OAM cell is to be received.
AIS tx	Time interval, in seconds, at which an AIS OAM cell is to be transmitted.
RDI rx	Time interval, in seconds, at which an RDI OAM cell is to be received.
RDI tx	Time interval, in seconds, at which an RDI OAM cell is to be transmitted.

NOTE There is a CLP0PPD threshold that is set by default as 87.5% of the buffer size and cannot be changed. (It is not displayed.)

14.3   Resetting a Network Module

This command is the software equivalent of removing and immediately re-inserting a network module while the switch is on. This allows you to reset the connections on a given network module without having physical access to the switch. Enter the following parameters:

```
myswitch::configuration module> reset <module>
```

This parameter is defined as follows:

Parameter	Description
module	The network module you want to reset.

NOTE All SVCs associated with this network module are torn down upon a reset and attempt to reconnect as needed. All PVCs associated with this network module are torn down and set up again upon a reset. Depending on the number of PVCs configured and the type of processor in your switch, it may take a short time for all of the PVCs to be re-established after a reset.

For example, if you want to reset network module 3B, enter the following:

```
myswitch::configuration module> reset 3b
```

The switch then cautions you that all connections will be torn down. To abort the reset, type **n** or press **<ENTER>**. To continue with the reset, type **y**.

```
Resetting the network module will destroy the
existing connections temporarily.
Reset the network module [n]? y
```

14.4   Displaying Network Module Information

This command lets you display general information about network modules that are currently installed in a switch fabric. Enter the following parameters:

```
myswitch::configuration module> show
Module Series Admin  Speed Ports Timing Rev. Serial#  ProductNumber
1A     C      up     100.0  6    no     1.0  N/A      NM-6/100MMSCC
1B     C      up     155.0  4    yes    1.1  N/A      NM-4/155MMSCC
1C     C2     up     1.536  4    yes    0.1  00000122 NMFR-4/DS1A
```

The fields in this display are defined as follows:

Field	Description
Module	Shows each network module currently installed in the switch fabric. 1 means that it is the switch fabric in slot 1. The letter shows the position of the network module in the switch fabric. E is the intra-fabric port to the other switch fabrics in an ASX-1000 or a TNX-1100 chassis.
Series	The hardware series of the network module. `C` is a Series C module. `LC` is a Series LC module. `C2` is a circuit emulation services (CES) module or a *FramePlus* module. `LE` is an interface group on an LE 155 or an LE 25 switch. `PC1` is a Series 1 logical network module on a port card in an ASX-4000. `D` is a Series D module.
Admin	`up` means this module is enabled. `down` means this module is disabled.
Speed	The speed in Mbps of the ports.
Ports	The number of ports on the network module.
Timing	`yes` means this network module supports distributed timing and `no` means that it does not.
Rev.	The hardware revision level of this network module.
Serial #[1]	The serial number of this network module.
Product Number	The FORE Systems product number for this module. `NM` is network module, `CE` is circuit emulation, `FR` is *FramePlus*, `BPB` is the ASX-1000 or TNX-1100 intra-fabric module.

[1.] This field is only available for Series D modules.

14.5   Configuring Traffic on a Network Module

These commands enable you to configure or to display information about the traffic on the network modules. To list the available subcommands, type **traffic ?** at the **module** level.

```
myswitch::configuration module> traffic ?
  c>                  lc>                 le>                 pc1>
  d>
```

The following is displayed on an LE 155 or LE 25:

```
myswitch::configuration module traffic> ?
  le>
```

The following is displayed on an ASX-4000:

```
myswitch::configuration module traffic> ?
  pc1>
```

14.5.1 Configuring Traffic on a Series C Network Module

These commands let you configure or display information about the traffic on Series C network modules. To list the available subcommands, type **c ?** at the **traffic** level.

```
myswitch::configuration module traffic> c ?
  epd             fifoblock       models          setmodel
  show
```

14.5.1.1 Setting Early Packet Discard on a Series C Network Module

This command lets you set a threshold value for AAL5 Early Packet Discard (EPD) on a specified network module. This is the static threshold (in terms of a specified percent of cells) at which EPD is activated. AAL5 frames that arrive when the shared buffer is over this threshold are discarded in whole. Enter the following parameters:

```
myswitch::configuration module traffic c> epd <module> <percent of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured for Early Packet Discard.
percent of cells	The AAL5 packet drop threshold to be set, in percent of cells. By default, this value is set to 90% of the shared buffer size.

14.5.1.2 Setting FIFO Blocking on a Series C Network Module

This command lets you set FIFO blocking on a specific network module. FIFO blocking enables buffering on the fabric when the network module buffers are full. Enter the following parameters:

```
myswitch::configuration module traffic c> fifoblock <module> (normal | enabled)
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured for FIFO blocking.
normal | enabled	**normal** means that the network module passes traffic normally. **enabled** means that FIFO blocking takes place on the network module when the buffers are full. The default is **normal**.

14.5.1.3 Displaying Traffic Models for a Series C Network Module

This command allows you to display the different types of traffic memory models on a network module. Enter the following parameters:

```
myswitch::configuration module traffic c> models
 Model   Memory Ucasts Mcasts   MOuts   Cells  Name
   1      32Kx48   4096    512    1024    2048  default
   2     128Kx48   6144    512    1024   12288  default
   3     128Kx48  11264   1024    2048   10240  more conns
   5     128Kx48   2048   2048   16384    8192  VP shaping
   6     128Kx48   3072    128    1024   13312  more cells
   7      32Kx48   2048    256    1024    2560  more cells
```

The fields in this display are defined as follows:

Field	Description
Model[1]	The shared memory model for this configuration. This is the number to enter for the `<model>` parameter when using the `conf module traffic c setmodel` command.
Memory	The size of this shared memory configuration, in bytes.
Ucasts[2]	The maximum number of unicast connections supported for this model. Half of these connections are reserved for low-priority connections (ABR, UBR) and half are reserved for high-priority connections (VBR, CBR). For example, if a model supports 4K unicast connections, it really supports 2K VBR or CBR connections and 2K ABR or UBR connections.
Mcasts	The number of input multicast connections supported from the switch fabric to the network module for this shared memory model. Half of these connections are reserved for low-priority connections (ABR, UBR) and half are reserved for high-priority connections (VBR, CBR). A multicast connection can have multiple outputs on the same network module. The number of outputs does not figure into this constraint.
MOuts	The number of output multicast connections supported from the network module to the link for this shared memory model. Any multicast connection can have multiple outputs on the same network module or the same port. If one multicast connection has four outputs, it requires 1 mcast and 4 mouts. Adding an output to a multicast connection can fail if all of the outputs are in use. Adding the multicast connection might not fail.
Cells	The total amount of cell buffering that is supported for this shared memory model.
Name	The identifier for this shared memory model.

[1.] Memory model 4 has been removed.

[2.] The Series C network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

14.5.1.4 Setting Traffic Models on a Series C Network Module

This command lets you select one of the traffic memory models for a specific network module. Enter the following parameters:

```
myswitch::configuration module traffic c> setmodel <module> <model>
```

These parameters are defined as follows:

Parameter	Description
module[1]	The network module to be configured.
model[2]	The predefined memory model to be used for this network module. The models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections. Enter the number found in the `Model` field of the **conf module traffic c models** command for the shared memory configuration that you want to use.

[1] The following Series C network modules can only use the following traffic memory models: NM-6/25UTPEC and NM-4/155UTP5EC can only use models 1 and 7; all other Series C network modules can only use models 2 - 6.

[2] The network module must be reset for this command to take effect.

14.5.1.5 Displaying Traffic on a Series C Network Module

This command enables you to display traffic model information about the Series C network modules. Enter the following parameters:

```
myswitch::configuration module traffic c> show
 Module  Memory  Model    EPD  FIFOblock
 4C      128Kx48 2        90%  enabled
 4D      128Kx48 2        90%  normal
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Memory	The hardware configuration this shared memory model supports, in units of 48-bit words (x48).
Model	The shared memory model used for this network module. See **conf module traffic c models** for more information.
EPD	The threshold for AAL5 Early Packet Discard on this network module in percent of the buffer size.
FIFOblock	Displays **enabled** if FIFO blocking is enabled on this network module. Otherwise, displays **normal**. The default is **normal**.

You can also display traffic model information about an individual Series C network module. Enter the following parameters:

```
myswitch::configuration module traffic c> show [<module>]
myswitch::configuration module traffic c> show 4C
 Module  Memory  Model    EPD  FIFOblock
 4C      128Kx48 2        90%  enabled
```

The fields in this display are defined in the same manner as those listed in the example above.

If no Series C network modules are installed, the following is displayed:

```
myswitch::configuration module traffic c> show
No Series-C traffic information is available
```

14.5.2 Configuring Traffic on a Series LC Network Module

These commands let you configure or display information about the traffic on Series LC network modules. To list the available subcommands, type **lc ?** at the **traffic** level.

```
myswitch::configuration module traffic> lc ?
  epd             efci           models          setmodel
  show
```

14.5.2.1 Setting EPD on a Series LC Network Module

This command lets you set a threshold value for AAL5 Early Packet Discard (EPD) on a Series LC network module. This is the static threshold (in terms of a specified percent of cells) at which EPD is activated. AAL5 packets that arrive when the shared buffer is over this threshold are discarded in whole. Enter the following parameters:

```
myswitch::configuration module traffic lc> epd <module> [UBR] <percent of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured for Early Packet Discard.
UBR	If UBR is entered, the AAL5 packet drop threshold is used for AAL5 UBR connections only. If UBR is not entered, the AAL5 packet drop threshold is used for all other AAL5 connections.
percent of cells	The AAL5 packet drop threshold to be set, in percent of cells. By default, this value is set to 90% of the shared buffer size.

14.5.2.2 Setting EFCI on a Series LC Network Module

This command lets you designate the cell buffer threshold over which Available Bit Rate (ABR) cells have their explicit forward congestion indicator (EFCI) bit set. When the EFCI bit is set, this signals congestion to downstream flow control mechanisms. Once this threshold is surpassed, EFCI continues to be set until the queue empties below the off threshold. Enter the following parameters:

```
myswitch::configuration module traffic lc> efci <module> (on | off)
<number of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The port on which the EFCI threshold will be set.
on | off[1]	**on** means the EFCI bit will be set when the threshold number is reached, signalling congestion. **off** means the EFCI bit will be cleared when the threshold number is reached, indicating no congestion.
number of cells	The number of cells for which the ABR[2] cells have the EFCI bit set or cleared. The default for the **on** threshold is 64. The default for the **off** threshold is 1.

[1.] The value for the **off** threshold must always be less than the value for the **on** threshold.

[2.] The Series LC network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

14.5.2.3 Displaying Traffic Models for a Series LC Network Module

This command allows you to display the different types of traffic memory models on a Series LC network module. Enter the following parameters:

```
myswitch::configuration module traffic lc> models
          Cell    Table
  Model   Memory  Memory Ucasts Mcasts  Cells Counters  Name
  1       256Kx64 32Kx32   6144    512  32768        0  default
  2       256Kx64 32Kx32   6144    512  16384        1  one counter
  3       256Kx64 32Kx32   6144    512   8192        2  two counters
  4       256Kx64 64Kx32   6144    512  32768        2  default
  5       512Kx64 64Kx32   6144    512  65536        2  default
  6       256Kx64 32Kx32   4096   1024  24576        1  more mcast
  7       256Kx64 64Kx32   4096   1024  32768        2  more mcast
  8       512Kx64 64Kx32   4096   1024  65536        2  more mcast
  9       256Kx64 32Kx32   6144    256  32768        0  8-port default
  10      256Kx64 32Kx32   4096    512  32768        0  8-port more mcast
  11      256Kx64 32Kx32   4096    512  24576        1  8-port one counter
  12      256Kx64 32Kx32   4096    512  14336        2  8-port two counters
```

The fields in this display are defined as follows:

Field	Description
Model[1]	The shared memory model for this configuration. This is the number to enter for the `<model>` parameter when using the **conf module traffic lc setmodel** command.
Cell Memory[2]	The hardware configuration this shared memory model supports, in units of 64-bit words (x64).
Table Memory	The hardware configuration this shared memory model supports, in units of 32-bit words (x32).
Ucasts	The maximum number of unicast connections supported by this shared memory configuration.
Mcasts[3]	The number of input multicast connections supported from the switch fabric to the network module by this shared memory configuration.
Cells	The total amount of cell buffering that is supported by this shared memory configuration.
Counters	The number of per-connection counters that is supported by this shared memory configuration. 0 means there are no per-connection counters. 1 means there are per-connection counters for cells transmitted and for cells lost. 2 means there are per-connection counters for cells transmitted, for cells lost, for intentional cells lost, and for unintentional cells lost.
Name	The identifier for this shared memory model.

[1.] For proper operation, all Series LC network modules in a switch must use memory models that support the same number of unicast connections. Therefore, all Series LC modules in a switch should either use models 1-5 OR models 6-8.

The following Series LC network modules can only use the following traffic memory models: OC-3 MM and UTP can only use models 1, 2, 3, and 6; OC-3 SM can only use models 4 and 7; and OC-12 MM and SM can only use models 5 and 8.

[2.] The Cell Memory and Table Memory sizes determine which memory models are appropriate for any given network module.

[3.] Only one multicast output is supported per port for each multicast connection.

14.5.2.4 Setting Traffic Models on a Series LC Network Module

This command lets you select one of the traffic memory models on a Series LC network module. Enter the following parameters:

```
myswitch::configuration module traffic lc> setmodel <module> <model>
```

These parameters are defined as follows:

Parameter	Description
module[1]	The network module to be configured.
model[2]	The predefined memory model to be used for this Series LC network module. The various models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections, and the number of per-connection counters. Enter the number found in the `Model` field of the **conf module traffic lc models** AMI command for the shared memory configuration that you want to use.

[1.] For proper operation, all Series LC network modules in a switch must use memory models that support the same number of unicast connections. Therefore, all Series LC modules in a switch should either use models 1-5 OR models 6-8.

The following Series LC network modules can only use the following traffic memory models: OC-3 MM and UTP can only use models 1, 2, 3, and 6; OC-3 SM can only use models 4 and 7; and OC-12 MM and SM can only use models 5 and 8.

[2.] The switch software must be restarted for this command to take effect.

14.5.2.5 Displaying Traffic on a Series LC Network Module

This command lets you display traffic model information about all of the Series LC network modules in a switch fabric. Enter the following parameters:

```
myswitch::configuration module traffic lc> show
            Cell     Table        CBR-VBR-ABR       UBR    EFCI   EFCI
  Module   Memory   Memory Model      EPD          EPD     On     Off
   1C     512Kx64   64Kx32    5       90%          90%     64      1
   1D     512Kx64   64Kx32    5       90%          90%     64      1
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Cell Memory	The hardware configuration this memory model supports (x64).
Table Memory	The hardware configuration this memory model supports (x32).
Model	The shared memory model used for this network module. See **conf module traffic lc models** for more information.
CBR-VBR-ABR EPD	The AAL5 packet drop threshold for CBR, VBR, and ABR traffic on this network module, in percent of the buffer size.
UBR EPD	The AAL5 packet drop threshold for UBR traffic on this network module, in percent of the buffer size.
EFCI On	The threshold value at which the EFCI will be set (turned on), signalling congestion, for ABR[1] traffic, in cells.
EFCI Off	The threshold value at which the EFCI will be cleared (turned off), indicating no congestion for ABR traffic, in cells.

[1.] The Series LC network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

You can also display traffic model information about an individual Series LC network module. Enter the following parameters:

```
myswitch::configuration module traffic lc> show [<module>]
myswitch::configuration module traffic lc> show 1D
            Cell     Table        CBR-VBR-ABR       UBR    EFCI   EFCI
  Module   Memory   Memory Model      EPD          EPD     On     Off
   1D     512Kx64   64Kx32    5       90%          90%     64      1
```

The fields in this display are defined in the same manner as those listed in the example above.

If no Series LC network modules are installed, the following is displayed:

```
myswitch::configuration module traffic lc> show
No Series-LC traffic information is available
```

Network Module Commands

14.5.3 Configuring Traffic on a Series LE Network Module

These commands let you configure or display information about the traffic on Series LE network modules. To list the available subcommands, type **le ?** at the **traffic** level.

```
myswitch::configuration module traffic> le ?
  epd              efci              models            setmodel
  show             vbrqueue
```

NOTE ➤ The **conf module traffic le** commands are only valid on an LE 155 or an LE 25 switch.

14.5.3.1 Setting EPD on a Series LE Network Module

This command lets you set a threshold value for AAL5 Early Packet Discard (EPD) on a Series LE network module. This is the static threshold (in terms of a specified percent of cells) at which EPD is activated. AAL5 packets that arrive when the shared buffer is over this threshold are discarded in whole. Enter the following parameters:

```
myswitch::configuration module traffic le> epd <module> [UBR] <percent of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured for Early Packet Discard.
UBR	On an LE 155, if **UBR** is entered, the AAL5 packet drop threshold is used for AAL5 UBR connections only. If **UBR** is not entered, the AAL5 packet drop threshold is used for all other AAL5 connections. On an LE 25, if **UBR** is entered, the AAL5 packet drop threshold is used for AAL5 UBR, ABR, and VBR connections by default. If the VBR traffic is moved into the high priority queue using the **conf module traffic le vbrqueue** command, then the AAL5 packet drop threshold is only used for AAL5 UBR and ABR connections. If **UBR** is not entered, the AAL5 packet drop threshold is used for AAL5 CBR connections by default. If the VBR traffic is moved into the high priority queue using the **conf module traffic le vbrqueue** command, then the AAL5 packet drop threshold is only used for AAL5 CBR and VBR connections.
percent of cells	The AAL5 packet drop threshold to be set, in percent of cells. By default, this value is set to 90% of the shared buffer size.

14.5.3.2 Setting EFCI on a Series LE Network Module

This command lets you designate the cell buffer threshold over which Available Bit Rate (ABR) cells have their explicit forward congestion indicator (EFCI) bit set. When the EFCI bit is set, this signals congestion to downstream flow control mechanisms. Once this threshold is surpassed, EFCI continues to be set until the queue empties below the `off` threshold. Enter the following parameters:

```
myswitch::configuration module traffic le> efci <module> (on | off)
<number of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The port on which the EFCI threshold will be set.
on | off[1]	**on** means the EFCI bit will be set when the threshold number is reached, signalling congestion. **off** means the EFCI bit will be cleared when the threshold number is reached, indicating no congestion.
number of cells	The number of cells for which the ABR[2] cells have the EFCI bit set or cleared. The default for the **on** threshold is 64. The default for the **off** threshold is 1.

[1.] The value for the **off** threshold must always be less than the value for the **on** threshold.

[2.] The Series LE network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

14.5.3.3 Displaying Traffic Models for a Series LE Network Module

This command allows you to display the different types of traffic memory models on a Series LE network module. Enter the following parameters:

```
myswitch::configuration module traffic le> models
          Cell   Table
 Model  Memory  Memory Ucasts Mcasts   Cells Counters  Name
   1    256Kx64  32Kx32   6144    512  32768        0  default
   2    256Kx64  32Kx32   6144    512  16384        1  one counter
   3    256Kx64  32Kx32   6144    512   8192        2  two counters
   4    256Kx64  64Kx32   6144    512  32768        2  default
   5    512Kx64  64Kx32   6144    512  65536        2  default
   6    256Kx64  32Kx32   4096   1024  24576        1  more mcast
   7    256Kx64  64Kx32   4096   1024  32768        2  more mcast
   8    512Kx64  64Kx32   4096   1024  65536        2  more mcast
   9    256Kx64  32Kx32   6144    256  32768        0  8-port default
  10    256Kx64  32Kx32   4096    512  32768        0  8-port more mcast
  11    256Kx64  32Kx32   4096    512  24576        1  8-port one counter
  12    256Kx64  32Kx32   4096    512  14336        2  8-port two counters
```

The fields in this display are defined as follows:

Field	Description
Model	The shared memory model for this configuration. This is the number to enter for the `<model>` parameter when using the **conf module traffic le setmodel** command.
Cell Memory[1]	The hardware configuration this shared memory model supports, in units of 64-bit words (x64).
Table Memory	The hardware configuration this shared memory model supports, in units of 32-bit words (x32).
Ucasts	The maximum number of unicast connections supported by this shared memory configuration.
Mcasts[2]	The number of input multicast connections supported from the switch fabric to the network module by this shared memory configuration.
Cells	The total amount of cell buffering that is supported by this shared memory configuration.
Counters	The number of per-connection counters that is supported by this shared memory configuration. 0 means there are no per-connection counters. 1 means there are per-connection counters for cells transmitted and for cells lost. 2 means there are per-connection counters for cells transmitted, for cells lost, for intentional cells lost, and for unintentional cells lost.
Name	The identifier for this shared memory model.

[1.] The `Cell` and `Table Memory` sizes determine which memory models are appropriate for a given network module.

[2.] Only one multicast output is supported per port for each multicast connection.

14.5.3.4 Setting Traffic Models on a Series LE Network Module

This command lets you select one of the traffic memory models on a Series LE network module. Enter the following parameters:

```
myswitch::configuration module traffic le> setmodel <module> <model>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured.
model[1]	The predefined memory model to be used for this Series LE network module. The various models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections, and the number of per-connection counters. Enter the number found in the `Model` field of the **conf module traffic le models** AMI command for the shared memory configuration that you want to use.

[1.] The switch software must be restarted for this command to take effect.

NOTE

For proper operation, all Series LE network modules in a switch must use memory models that support the same number of unicast connections. LE 155 switches can only use the following traffic memory models: 1, 2, 3, and 6. Therefore, all network modules in LE 155 switches should either use models 1-3 <u>or</u> 6.

LE 25 switches can only use the following traffic memory models: 9, 10, 11, and 12. Therefore, all network modules, except the Port Expansion Module (PEM), in LE 25 switches should either use model 9 <u>or</u> models 10-12. The network modules in LE 25 switches default to model 9, and the PEM defaults to model 1.

Network Module Commands

14.5.3.5 Displaying Traffic on a Series LE Network Module

This command lets you display traffic model information on the Series LE network modules.
The following is displayed on an LE 155 switch:

```
myswitch::configuration module traffic le> show
          Cell   Table           CBR-VBR-ABR   UBR   EFCI   EFCI
Module  Memory  Memory  Model           EPD   EPD    On    Off
 1B     256Kx64  32Kx32  2             90%   90%    64      1
 1C     256Kx64  32Kx32  2             90%   90%    64      1
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Cell Memory	The size of this shared memory cell RAM configuration (x64).
Table Memory	The size of this shared memory table RAM configuration (x32).
Model	The shared memory model used for this network module. See `conf module traffic le models` for more information.
CBR-VBR-ABR EPD	The AAL5 packet drop threshold for CBR, VBR, and ABR traffic on this network module, in percent of the buffer size.
UBR EPD	The AAL5 packet drop threshold for UBR traffic on this network module, in percent of the buffer size.
EFCI On	The threshold value at which the EFCI will be set (turned on), signalling congestion, for ABR[1] traffic, in cells.
EFCI Off	The threshold value at which the EFCI will be cleared (turned off), indicating no congestion for ABR traffic, in cells.

[1.] The Series LE network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

You can also display traffic model information about an individual Series LE network module.
The following is displayed on an LE 155 switch:

```
myswitch::configuration module traffic le> show [<module>] [vbrqueue]
myswitch::configuration module traffic le> show 1B
          Cell   Table           CBR-VBR-ABR   UBR   EFCI   EFCI
Module  Memory  Memory  Model           EPD   EPD    On    Off
 1B     256Kx64  32Kx32  2             90%   90%    64      1
```

The fields in this display are defined in the same manner as those listed in the example above.

Something similar to the following is displayed on an LE 25 switch. In this example, the VBR traffic on module 1B has been placed in the high priority queue using `conf module traffic le vbrqueue` *<module>* `rt`. The VBR traffic on modules 1A and 1C is in the low priority queue. Module 1D is the Port Expansion Module (PEM).

```
myswitch::configuration module traffic le> show
           Cell     Table                CBR  VBR-ABR-UBR   EFCI    EFCI
 Module   Memory   Memory Model          EPD          EPD    On     Off
 1A       256Kx64  32Kx32     9          90%          90%    64       1
           Cell     Table             CBR-VBR      ABR-UBR   EFCI    EFCI
 Module   Memory   Memory Model          EPD          EPD    On     Off
 1B       256Kx64  32Kx32     9          90%          90%    64       1
           Cell     Table                CBR  VBR-ABR-UBR   EFCI    EFCI
 Module   Memory   Memory Model          EPD          EPD    On     Off
 1C       256Kx64  32Kx32     9          90%          90%    64       1
           Cell     Table         CBR-VBR-ABR          UBR   EFCI    EFCI
 Module   Memory   Memory Model          EPD          EPD    On     Off
 1D       256Kx64  32Kx32     1          90%          90%    64       1
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Cell Memory	The size of this shared memory cell RAM configuration (x64).
Table Memory	The size of this shared memory table RAM configuration (x32).
Model	The shared memory model used for this network module. See `conf module traffic le models` for more information.
CBR EPD	The AAL5 packet drop threshold for CBR traffic only on this network module, in percent of the buffer size, when the VBR traffic is in the low priority queue (the default condition).
VBR-ABR-UBR EPD	The AAL5 packet drop threshold for VBR, ABR, and UBR traffic on this network module, in percent of the buffer size. This field applies to modules that have VBR traffic placed in the low priority queue (the default condition).
CBR-VBR EPD	The AAL5 packet drop threshold for CBR and VBR traffic on this network module, in percent of the buffer size, when the VBR traffic is in the high priority queue.
ABR-UBR EPD	The AAL5 packet drop threshold for ABR and UBR traffic on this network module, in percent of the buffer size. This field applies to modules that have VBR traffic placed in the high priority queue.
CBR-VBR-ABR EPD	The AAL5 packet drop threshold for CBR, VBR, and ABR traffic on this network module, in percent of the buffer size. This field only applies to the PEM (module 1D).
UBR EPD	The AAL5 packet drop threshold for UBR traffic on this network module, in percent of the buffer size. This field only applies to the PEM (module 1D).

Field	Description
EFCI On	The threshold value at which the EFCI will be set (turned on), signalling congestion, for ABR[1] traffic, in cells.
EFCI Off	The threshold value at which the EFCI will be cleared (turned off), indicating no congestion for ABR traffic, in cells.

[1.] The Series LE network modules offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

You can display traffic model information about an individual Series LE network module on an LE 25 switch:

```
myswitch::configuration module traffic le> show [<module>] [vbrqueue]
myswitch::configuration module traffic le> show 1B
           Cell     Table              CBR-VBR     ABR-UBR   EFCI   EFCI
  Module   Memory   Memory Model          EPD         EPD    On     Off
   1B      256Kx64  32Kx32    9           90%         90%    64      1
```

The fields in this display are defined in the same manner as those listed in the example above.

You can display VBR queue information on an LE 25 switch as follows:

```
myswitch::configuration module traffic le> show [<module>] [vbrqueue]
myswitch::configuration module traffic le> show vbrqueue
Module    VBR-Queue
1A        nrt
1B        rt
1C        nrt
1D        N/A
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
VBR-Queue	The queue that is servicing the VBR traffic. **rt** is the high priority queue and **nrt** is the low priority queue. This field does not apply to the PEM (module 1D).

If no Series LE network modules are installed (this is not an LE 155 or an LE 25 switch), then the following is displayed:

```
myswitch::configuration module traffic le> show
No Series-LE traffic information is available
```

14.5.3.6 Configuring the VBR Traffic on a Series LE Network Module

This command lets you treat VBR traffic as either nrt-VBR traffic (the traffic is placed in the low priority queue) or as rt-VBR traffic (the traffic is placed in the high priority queue) on an network module in an LE 25 switch.

There are two queues for servicing traffic: high priority and low priority. By default, the high priority queue services CBR traffic, and the low priority queue services VBR, ABR, and UBR traffic.

The default configuration is ideal if you want to connect your voice network to an LE 25. However, if your network primarily uses data, you may want to move the VBR traffic into the high priority queue. In this way, the VBR traffic is not affected by the lower priority traffic.

> **NOTE**   This command only applies to an LE 25 switch. It does not apply to an LE 155 switch.

Enter the following parameters:

```
myswitch::configuration module traffic le> vbrqueue <module> (rt | nrt)
```

These parameters are defined as follows:

Parameter	Description
module	The network module in an LE 25 switch that is to be configured. Can only be 1A, 1B, or 1C. This command does not apply to the Port Expansion Module (PEM) (1D).
(rt | nrt)	**rt** places the VBR traffic in the high priority queue. **nrt** places the VBR traffic in the low priority queue. The default is **nrt**.

When you use this command, the switch asks you to reset the network module. For example:

```
myswitch::configuration module traffic le> vbrqueue 1a rt
Changing the VBR-queue requires a reset of the network module.
Continue with the change [n]? y
```

Entering **n** or pressing **<ENTER>** aborts the command. Entering **y** resets the network module and moves the VBR traffic to the designated queue.

14.5.4 Configuring Traffic on a Series 1 Port Card

On an ASX-4000, each port card contains two logical network modules. These commands let you configure or display information about the traffic on the logical network modules on Series 1 port cards. To list the available subcommands, type **?** at the `pc1` level.

```
myswitch::configuration module traffic pc1> ?
  epd              efci              models            setmodel
  show
```

14.5.4.1 Setting EPD on a Series 1 Port Card

This command lets you set a threshold value for AAL5 Early Packet Discard (EPD) on the logical network modules on Series 1 port cards. This is the static threshold (in terms of a specified percent of cells) at which EPD is activated. AAL5 packets that arrive when the shared buffer is over this threshold are discarded in whole. Enter the following parameters:

```
myswitch::configuration module traffic pc1> epd <module> [UBR] <percent of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The logical network module to be configured for Early Packet Discard.
UBR	If **UBR** is entered, the AAL5 packet drop threshold is used for AAL5 UBR connections only. If **UBR** is not entered, the AAL5 packet drop threshold is used for all other AAL5 connections.
percent of cells	The AAL5 packet drop threshold to be set, in percent of cells. By default, this value is set to 90% of the shared buffer size.

> **NOTE** This command is not available on Series 1 OC-48c port cards.

14.5.4.2 Setting EFCI on a Series 1 Port Card

This command lets you designate the cell buffer threshold over which Available Bit Rate (ABR) cells have their explicit forward congestion indicator (EFCI) bit set. When the EFCI bit is set, this signals congestion to downstream flow control mechanisms. Once this threshold is surpassed, EFCI continues to be set until the queue empties below the off threshold. Enter the following parameters:

```
myswitch::configuration module traffic pc1> efci <module> (on | off)
<number of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The logical network module on which the EFCI threshold will be set.
on | off[1]	**on** means the EFCI bit will be set when the threshold number is reached, signalling congestion. **off** means the EFCI bit will be cleared when the threshold number is reached, indicating no congestion.
number of cells	The number of cells for which the ABR[2] cells will have the EFCI bit set or cleared. The default for the **on** threshold is 64. The default for the **off** threshold is 1.

[1.] The value for the **off** threshold must always be less than the value for the **on** threshold.

[2.] The Series 1 port cards offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

NOTE ▶ This command is not available on Series 1 OC-48c port cards.

14.5.4.3 Displaying Traffic Models for a Series 1 Port Card

This command allows you to display the different types of traffic memory models for a Series 1 port card. Enter the following parameters:

```
myswitch::configuration module traffic pc1> models
          Cell    Table
 Model   Memory  Memory Ucasts Mcasts  Cells Counters  Name
   1    1024Kx64 128Kx32  24576    512 131072        0  mcast512 -- default
   2    1024Kx64 128Kx32  24576    512  65536        1  mcast512 -- one counter
   3    1024Kx64 128Kx32  24576    512  32768        2  mcast512 -- two counters
   4    1024Kx64 256Kx32  24576    512 131072        2  mcast512 -- default
   5    2048Kx64 256Kx32  24576    512 262144        2  mcast512 -- default
   6    1024Kx64 128Kx32  16384   1024  98304        1  mcast1K
   7    1024Kx64 256Kx32  16384   1024 131072        2  mcast1K
   8    2048Kx64 256Kx32  16384   1024 262144        2  mcast1K
   9    2048Kx64 256Kx32  16384   2048 229376        2  mcast2K
```

The fields in this display are defined as follows:

Field	Description
Model[1]	The shared memory model for this configuration. This is the number to enter for the `<model>` parameter when using the `conf module traffic pc1 setmodel` command.
Cell Memory[2]	The hardware configuration this shared memory model supports, in 64-bit words (x64).
Table Memory	The hardware configuration this shared memory model supports, in 32-bit words (x32).
Ucasts	The maximum number of unicast connections supported by this shared memory configuration.
Mcasts[3]	The number of input multicast connections supported from the switch fabric to the logical network module by this shared memory configuration.
Cells	The total amount of cell buffering that is supported by this shared memory configuration.
Counters	The number of per-connection counters that is supported by this shared memory configuration. `0` means there are no per-connection counters. `1` means there are per-connection counters for cells transmitted and for cells lost. `2` means there are per-connection counters for cells transmitted, for cells lost, for intentional cells lost, and for unintentional cells lost.
Name	The identifier for this shared memory model.

[1.] For proper operation, all Series 1 port cards in a fabric must use memory models that support the same number of unicast connections. Additionally, their multicast spaces (indicated by their `Name`) must be compatible with what you have configured under `conf board setmcastspace`.
The Series 1 OC-12c port cards can only use the following traffic memory models: 5, 8, and 9. This command does not apply to Series 1 OC-48c port cards.

[2.] The `Cell` and `Table Memory` sizes determine which memory models are appropriate for a given logical network module.

[3.] Only one multicast output is supported per port for each multicast connection.

14.5.4.4 Setting Traffic Models on a Series 1 Port Card

This command lets you select one of the traffic memory models on a Series 1 port card. Enter the following parameters:

```
myswitch::configuration module traffic pc1> setmodel <module> <model>
```

These parameters are defined as follows:

Parameter	Description
module[1]	The logical network module to be configured.
model[2]	The predefined memory model to be used for this logical network module on a Series 1 port card. The various models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections, and the number of per-connection counters. Enter the number found in the Model field of the **conf module traffic pc1 models** command for the shared memory configuration that you want to use.

[1.] For proper operation, all Series 1 port cards in a fabric must use memory models that support the same number of unicast connections. Therefore, all Series 1 port cards in a fabric should either use model 5 <u>or</u> models 8-9.

[2.] The switch software must be restarted for this command to take effect.

For example, if you attempt to change the memory model to model 5, but that it is the current model, the switch returns a message as follows:

```
myswitch::configuration module traffic pc1> setmodel 1c 5
NOTICE: There is no change in the memory model.
No action will be taken.
```

> **NOTE** The model that is used must be compatible with the multicast space configuration under **conf board setmcastspace**. See Section 5.2 in this manual for more information about this command.

For example, if you want to change the memory model to model 8, but the multicast space configuration does not support that model, the switch returns an error message as follows:

```
myswitch::configuration module traffic pc1> setmodel 1c 8
WARNING: For correct operation, memory models for
all Series-PC1 network modules on this board as well
as this board's multicast space should be compatible.
Multicast space for a board can be configured via
<configure board setmcastspace> command.
```

```
Changing the configured model on a Series-PC1
network module will take effect after restarting
the switch software.

Set memory model to 8 [y]? y

?ERROR: jalapeno_set_config_id: Memory model 8 is not compatible with board 1's
multicast space. Please set this board's multicast space to a compatible model
before changing the memory model to 8 for network module 1C
```

You need to change the multicast space configuration to a compatible model. Look in the Name field under **conf module traffic pc1 models**. The Name for model 8 is mcast1k. So, you need to change the multicast space configuration to index 2 (mcast1k) and restart the switch as follows:

```
myswitch::configuration module traffic pc1> conf board setmcastspace 1 2

WARNING: For correct operation, all OC3/OC12 Series PC1
network modules on  board 1 must  use  memory  models
that  support  mcast1k. OC3/OC12  Series  PC1  port
cards may be autoconfigured to a new memory model
that supports this multicast space.

Changing the multicast space model on a board will take
effect after restarting the switch software.

Set multicast space index to 2 [y]? y

Restart the switch software [n]? y
Connection closed by foreign host.
```

After you log in again, you can see that the memory model has been automatically updated to a compatible model:

```
myswitch::configuration module traffic pc1> show 1c
          Cell    Table           CBR-VBR-ABR   UBR   EFCI   EFCI
 Module  Memory  Memory Model           EPD   EPD    On    Off
 1C     2048Kx64 256Kx32 8             90%   90%    64     1
```

14.5.4.5 Displaying Traffic on a Series 1 Port Card

This command lets you display traffic model information on the Series 1 port cards. The information is displayed for each module in the order in which they are installed in the switch. In this example, 2A and 2B are logical network modules in an OC-12c port card and 2D is a logical network module in an OC-48c port card. Enter the following parameters:

```
myswitch::configuration module traffic pc1> show
          Cell     Table           CBR-VBR-ABR   UBR    EFCI   EFCI
  Module  Memory   Memory  Model          EPD    EPD    On     Off
  2A      2048Kx64 256Kx32 5              90%    90%    64     1
  2B      2048Kx64 256Kx32 5              90%    90%    64     1
          Ucast conns        Mcast conns
  Module  Current    Max    Current    Max
  2D            5  63488          0   2048
```

The fields in this display are defined as follows:

Field	Description
Module	The logical network module that has been configured.
Cell Memory	The hardware configuration this memory model supports (x64).
Table Memory	The hardware configuration this memory model supports (x32).
Model	The shared memory model used for this logical network module. See `conf module traffic pc1 models` for more information.
CBR-VBR-ABR EPD	The AAL5 packet drop threshold for CBR, VBR, and ABR traffic on this logical network module, in percent of the buffer size.
UBR EPD	The AAL5 packet drop threshold for UBR traffic on this logical network module, in percent of the buffer size.
EFCI On	The threshold value at which the EFCI will be set (turned on), signalling congestion, for ABR[1] traffic, in cells.
EFCI Off	The threshold value at which the EFCI will be cleared (turned off), indicating no congestion for ABR traffic, in cells.
Current Ucast conns	The current number of active unicast connections on this logical network module.
Max Ucast conns	The total number of unicast connections that this logical network module can support.
Current Mcast conns	The current number of active multicast connections on this logical network module.
Max Mcast conns	The total number of multicast connections that this logical network module can support.

[1.] The Series 1 port cards offer an ABR-ready ATM interface. A future *ForeThought* release will support ABR QoS operation on these interfaces. ABR options in the current release will apply in that *ForeThought* release.

You can also display traffic model information about an individual logical network module on a Series 1 port card. Enter the following parameters:

```
myswitch::configuration module traffic pc1> show [<module>]
myswitch::configuration module traffic pc1> show 2a
            Cell    Table        CBR-VBR-ABR    UBR   EFCI   EFCI
 Module   Memory   Memory  Model        EPD    EPD    On    Off
 2A       2048Kx64 256Kx32  5           90%    90%    64     1
```

The fields in this display are defined in the same manner as those listed in the example above.

If no Series 1 port cards are installed, the following is displayed:

```
myswitch::configuration module traffic pc1> show
No Series-PC1 traffic information is available
```

14.5.5 Configuring Traffic on a Series D Network Module

These commands let you configure or display information about the traffic on Series D network modules. To list the available subcommands, type **?** at the **d** level.

```
myswitch::configuration module traffic d> ?
  aal5pktcount    altclpthresh    epd            efci
  models          setmodel        show           vcclpthresh
```

14.5.5.1 Enabling/Disabling the AAL5 Transmitted PDU Counter

This command lets you enable or disable the per-connection AAL5 transmitted Protocol Data Unit (PDU) counter on a per network module basis on a Series D network module. Enter the following parameters:

```
myswitch::configuration module traffic d> aal5pktcount <module>
(enable | disable)
```

These parameters are defined as follows:

Parameter	Description
module[1]	The network module to be configured for counting AAL5 packets.
enable | disable	**enable** means the network module will count the number of transmitted AAL5 PDUs. **disable** means the network module will not count the number of transmitted AAL5 PDUs. The default is **disable**.

[1.] You must reset the network module once you have entered this command.

14.5.5.2 Configuring Alternate CLP Thresholds

This command allows you to set the per-connection alternate CLP thresholds on a per network module basis. Enter the following:

```
myswitch::configuration module traffic d> altclpthresh <module> (clp1 | clp01)
<threshold>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured with a new CLP threshold.
clp1	The threshold is being set for CLP=1 cells. Applying a CLP=1 threshold means that when the current cell count for a connection is greater than this threshold, cells that have a CLP=1 are dropped. This threshold must be less than the CLP=0+1 threshold.

Parameter	Description
clp01	The threshold is being set for CLP=0+1 cells. Applying a CLP=0+1 threshold means that when the current cell count for a connection is greater than this threshold, cells are dropped, regardless of their CLP bit. This threshold must be greater than the CLP=1 threshold.
threshold	The number of cells in the per-connection buffer at which the connection drops CLP=1 or CLP=0+1 cells. These thresholds can only be adjusted in increments of multiples of 16 cells (e.g., 16, 32, 48). For CLP=1, the minimum value is 0, the maximum is the CLP=0+1 threshold - 16 cells, and the default is 256 cells. For CLP=0+1, the minimum value is the CLP=1 threshold + 16 cells, the maximum is 4,080 cells, and the default is 4,080 cells.

14.5.5.3 Setting EPD on a Series D Network Module

This command lets you set a threshold value for AAL5 Early Packet Discard (EPD) for UBR traffic and a different threshold value for AAL5 EPD for CBR, VBR, and ABR traffic on a Series D network module. This works differently than it does for other network modules. On a Series D network module, this is the static threshold (in terms of a specified percent of cells) at which EPD is activated. AAL5 packets are discarded in whole if they arrive when the total cell buffering is over this threshold and when the corresponding dedicated queue is full. Enter the following:

```
myswitch::configuration module traffic d> epd <module> [ubr] <percent of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured for Early Packet Discard.
ubr	If ubr is entered, the AAL5 packet drop threshold is used for AAL5 UBR connections only. If ubr is not entered, the AAL5 packet drop threshold is used for CBR, VBR, and ABR AAL5 connections.
percent of cells	The AAL5 packet drop threshold to be set, in percent of cells. The default value is set to 90% of the total cell buffer size for both the UBR and non-UBR thresholds.

14.5.5.4 Setting EFCI on a Series D Network Module

This command lets you designate the cell buffer threshold over which all cells have their explicit forward congestion indicator (EFCI) bit set. When the EFCI bit is set, this signals congestion to downstream flow control mechanisms. Once this threshold is surpassed, EFCI continues to be set until the queue empties below the off threshold. Enter the following parameters:

```
myswitch::configuration module traffic d> efci <module> (on | off)
<number of cells>
```

These parameters are defined as follows:

Parameter	Description
module	The network module on which the EFCI threshold will be set.
on | off[1]	**on** means the EFCI bit will be set when the threshold number is reached, signalling congestion. **off** means the EFCI bit will be cleared when the threshold number is reached, indicating no congestion.
number of cells	The number of cells for which the cells have the EFCI bit set or cleared. This number must be entered in increments of 32. The default for the **on** threshold is 256. The default for the **off** threshold is 192.

[1.] The value for the **off** threshold must always be less than or equal to the value for the **on** threshold.

14.5.5.5 Displaying Traffic Models for a Series D Network Module

This command allows you to display the different types of traffic memory models on a Series D network module. Enter the following parameters:

```
myswitch::configuration module traffic d> models
          Cell   Table
 Model  Memory  Memory Ucasts Mcasts  MOuts  Cells Counters  Name
   1    256Kx32 256Kx32  12288   1024   4096  16352        8  16K cells OC-12
   2    256Kx32 256Kx32  12288   1024   4096  16352        8  16K cells OC-3
   3    256Kx32 256Kx32  10240    512   4096  16352        8  16K cells DSn/En
   4    512Kx32 256Kx32  12288   1024   4096  32736        8  32K cells OC-12
   5    512Kx32 256Kx32  12288   1024   4096  32736        8  32K cells OC-3
   6    512Kx32 256Kx32  10240    512   4096  32736        8  32K cells DSn/En
```

The fields in this display are defined as follows:

Field	Description
Model[1]	The shared memory model for this configuration. This is the number to enter for the *<model>* parameter when using the **conf module traffic d setmodel** command.
Cell Memory[2]	The hardware configuration this shared memory model supports, in units of 32-bit words (x32).
Table Memory	The hardware configuration this shared memory model supports, in units of 32-bit words (x32).
Ucasts	The maximum number of unicast connections supported by this shared memory configuration.
Mcasts	The number of input multicast connections supported from the switch fabric to the network module by this shared memory configuration.
MOuts	The number of output multicast connections supported from the network module to the link for this shared memory model. Any multicast connection can have multiple outputs on the same network module or the same port. If one multicast connection has four outputs, it requires 1 mcast and 4 mouts.
Cells	The total amount of cell buffering that is supported by this shared memory configuration.
Counters[3]	The number of per-connection counters that is supported by this shared memory configuration. The counters are: 1 - Transmit (CLP=0) counter, 2 - Transmit (CLP=1) counter, 3 - EPD counter, 4 - Unintentional cell loss counter, 5 - CLP=1 threshold counter, 6 - CLP=0+1 threshold counter, 7 - Intentional PPD counter, and 8 - Transmit packet counter.
Name	The identifier for this shared memory model.

[1.] The following Series D network modules can only use the following traffic memory models: OC-12 network modules can only use model 4; OC-3 network modules can only use model 5; and DS3, E3, DS1, and E1 network modules can only use model 6. Models 1, 2, and 3 are not used.

[2.] The Cell and Table Memory sizes determine which memory models are appropriate for a given network module.

[3.] If the module packet counter is enabled under **conf module traffic d aal5pktcount**, the Transmit CLP=0 and CLP=1 counters are combined into a single Transmit CLP0+1 count.

14.5.5.6 Setting Traffic Models on a Series D Network Module

This command lets you select one of the traffic memory models on a Series D network module. Enter the following parameters:

```
myswitch::configuration module traffic d> setmodel <module> <model>
```

These parameters are defined as follows:

Parameter	Description
module[1]	The network module to be configured.
model[2]	The predefined memory model to be used for this Series D network module. The various models make different trade-offs between the number of cell buffers, and the number of unicast and multicast connections, and the number of per-connection counters. Enter the number found in the `Model` field of the **conf module traffic d models** AMI command for the shared memory configuration that you want to use.

[1.] The following Series D network modules can only use the following traffic memory models: OC-12 network modules can only use model 4; OC-3 network modules can only use model 5; and DS3, E3, DS1, and E1 network modules can only use model 6. Models 1, 2, and 3 are not used.

[2.] The network module must be reset for this command to take effect.

NOTE ▶ Since there is currently only one valid model for each type of network module, the software will not allow you to change the memory model for a Series D network module.

Network Module Commands

14.5.5.7 Displaying Traffic on a Series D Network Module

This command lets you display traffic model information about the Series D network modules. Enter the following parameters:

```
myswitch::configuration module traffic d> show
          Cell          CBR-VBR-ABR   UBR    EFCI   EFCI    AAL5   AltVC Thresholds
Module  Memory Model          EPD    EPD     On    Off PktCnt    CLP0+1     CLP1
1A     256Kx32              14717  14717    256    192 disable     4080      256
1B     256Kx32    1         14717  14717    256    192 disable     4080      256
1C     256Kx32    3         14717  14717    256    192 disable     4080      256
1D     256Kx32    3         14717  14717    256    192 disable     4080      256
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Cell Memory	The size of this shared memory cell RAM configuration (x32).
Model	The shared memory model used for this network module. See **conf module traffic d models** for more information.
CBR-VBR-ABR EPD	The AAL5 packet drop threshold for CBR, VBR, and ABR traffic on this network module, in cells.
UBR EPD	The AAL5 packet drop threshold for UBR traffic on this network module, in cells.
EFCI On	The threshold value at which the EFCI bit will be set (turned on), signalling congestion, for all traffic, in cells.
EFCI Off	The threshold value at which the EFCI bit will be cleared (turned off), indicating no congestion for all traffic, in cells.
AAL5 PktCnt	**enable** means the network module counts the number of transmitted AAL5 PDUs. **disable** means the network module does not count the number of transmitted AAL5 PDUs. The default is **disable**.
AltVC Threshold CLP0+1	The alternate threshold at which cells are dropped, regardless of their CLP bit, when the current cell count for a connection is greater than this threshold. This threshold is used instead of the VC CLP0+1 threshold if the **conf port traffic d altclpconfig** command is applied or if it is enabled using the **-AltCLP** option under **conf upc new**.
AltVC Threshold CLP1	The alternate threshold at which cells that are CLP=1 are dropped when the current cell count for a connection is greater than this threshold. This threshold is used instead of the VC CLP1 threshold if the **conf port traffic d altclpconfig** command is applied or if it is enabled using the **-AltCLP** option under **conf upc new**.

You can also display traffic model information about an individual Series D network module. Enter the following parameters:

```
myswitch::configuration module traffic d> show [<module>] [perclass]
myswitch::configuration module traffic d> show 1A
          Cell       CBR-VBR-ABR   UBR   EFCI   EFCI   AAL5   AltVC Thresholds
 Module  Memory Model       EPD   EPD    On    Off  PktCnt   CLP0+1    CLP1
 1A      256Kx32   2       14717 14717   256    192 disable    4080     256
```

The fields in this display are defined in the same manner as those listed previously.

You can also display traffic information about the Series D network modules on a per-class basis. Enter the following parameters:

```
myswitch::configuration module traffic d> show [<module>] [perclass]
myswitch::configuration module traffic d> show perclass
                  VC Thresholds
 Module  Priority CLP0+1    CLP1
 1A           CBR   3072    2048
 1A           VBR   3072    2048
 1A           UBR   3072    2048
 1A           ABR   3072    2048
 1B           CBR   3072    2048
 1B           VBR   3072    2048
 1B           UBR   3072    2048
 1B           ABR   3072    2048
 1C           CBR   3072    2048
 1C           VBR   3072    2048
 1C           UBR   3072    2048
 1C           ABR   3072    2048
 1D           CBR   3072    2048
 1D           VBR   3072    2048
 1D           UBR   3072    2048
 1D           ABR   3072    2048
```

The fields in this display are defined as follows:

Field	Description
Module	The network module that has been configured.
Priority	The class of service to which these thresholds apply.
VC Thresholds CLP0+1	The threshold at which cells are dropped, regardless of their CLP bit, when the current cell count for a connection is greater than this threshold for this traffic type.
VC Thresholds CLP1	The threshold at which cells that are CLP=1 are dropped when the current cell count for a connection is greater than this threshold for this traffic type.

If no Series D network modules are installed, then the following is displayed:

```
myswitch::configuration module traffic d> show
No Series-D traffic information is available
```

14.5.5.8 Configuring the VC CLP Threshold on a Series D Network Module

This command lets you configure the per-connection CLP threshold for all VCs on a per-class basis. Enter the following parameters:

```
myswitch::configuration module traffic d> vcclpthresh <module>
(cbr | vbr | abr | ubr) (clp1 | clp01) <threshold>
```

These parameters are defined as follows:

Parameter	Description
module	The network module to be configured.
cbr	vbr	abr	ubr	The class of service for which this threshold applies.
clp1	The threshold is being set for CLP=1 cells. Applying a CLP=1 threshold means that when the current cell count for a connection is greater than this threshold, cells that have a CLP=1 are dropped. This threshold must be less than the CLP=0+1 threshold.
clp01	The threshold is being set for CLP=0+1 cells. Applying a CLP=0+1 threshold means that when the current cell count for a connection is greater than this threshold, cells are dropped, regardless of their CLP bit. This threshold must be greater than the CLP=1 threshold.
threshold	The number of cells in the buffer at which the specified traffic type drops CLP=1 or CLP=0+1 cells. This threshold can only be adjusted in increments of multiples of 16 cells (e.g., 16, 32, 48). For CLP=1, the minimum value is 0, the maximum is the CLP=0+1 threshold - 16 cells, and the default is 256 cells. For CLP=0+1, the minimum value is the CLP=1 threshold + 16 cells, the maximum is 4,080 cells, and the default is 4,080 cells.

CHAPTER 15 NSAP Configuration Commands

These commands allow you to configure NSAP prefixes; to display ILMI registered NSAP addresses; to configure NSAP-to-E.164 address mappings; and to configure a table of LECS addresses that switches can query when attached LECs ask for the registered LECS address. To list the available commands, type **?** at the **nsap** level.

```
myswitch::configuration nsap> ?
  prefix>          ilmi>          e164>          registry>
```

15.1 NSAP Prefix Configuration Commands

These commands enable you to delete an NSAP prefix, create an NSAP prefix, and display NSAP prefix information. You can display the list of available subcommands by typing **?** at the **prefix** level.

```
myswitch::configuration nsap prefix> ?
  delete           new             show
```

15.1.1 Deleting an NSAP Prefix

This command lets you remove an existing NSAP prefix. Enter the following parameters:

```
myswitch::configuration nsap prefix> delete <port> <vpi> <prefix>
```

15.1.2 Creating an NSAP Prefix

This command lets you create an NSAP prefix. Enter the following parameters:

```
myswitch::configuration nsap prefix> new <port> <vpi> <prefix>
```

The parameters for delete and new are defined as follows:

Parameter	Description
port	The port number for this NSAP prefix.
vpi	The virtual path number for this NSAP prefix.
prefix[1]	The NSAP prefix for this entry.

[1.] Because multiple prefixes are not supported on the same port, delete the old prefix before creating a new one.

15.1.3 Displaying NSAP Prefixes

This command enables you to display the current list of NSAP prefixes. Enter the following parameters:

```
myswitch::configuration nsap prefix> show
Port      VPI NSAP-Prefix
1A1         0 0x47.0005.80.ffe100.0000.f124.00de
1A2         0 0x47.0005.80.ffe100.0000.f124.00de
1B1         0 0x47.0005.80.ffe100.0000.f124.00de
1B2         0 0x47.0005.80.ffe100.0000.f124.00de
1CTL        0 0x47.0005.80.ffe100.0000.f124.00de
```

The fields in this display are defined as follows:

Field	Description
Port	The port number for this NSAP prefix.
VPI	The virtual path number on which this NSAP prefix exists.
NSAP-Prefix	The NSAP prefix for this entry.

You can also display information about an individual port; or port and vpi; or port, vpi, and prefix as follows:

```
myswitch::configuration nsap prefix> show [<port> [<vpi> [<prefix>]]]
myswitch::configuration nsap prefix> show 1a1
Port      VPI NSAP-Prefix
1A1         0 0x47.0005.80.ffe100.0000.f124.00de


myswitch::configuration nsap prefix> show 1a1 0
Port      VPI NSAP-Prefix
1A1         0 0x47.0005.80.ffe100.0000.f124.00de


myswitch::configuration nsap prefix> show 1a1 0
0x47.0005.80.ffe100.0000.f124.00de
Port      VPI NSAP-Prefix
1A1         0 0x47.0005.80.ffe100.0000.f124.00de
```

The fields in these displays are defined in the same manner as those listed previously.

If you have not configured any NSAP prefixes, then the following message is displayed:

```
myswitch::configuration nsap prefix> show
No user configured NSAP prefix information is available
```

15.2 NSAP ILMI Configuration Command

This command lets you display the NSAP addresses of all of the ports on a switch fabric that have been registered via ILMI. ILMI address registration occurs between the switch and host. The switch sends the host its 13-byte NSAP prefix. If the host accepts the prefix, the host builds its own NSAP address by appending its 7-byte host specific part. The host returns the complete 20-byte NSAP address to the switch. If the switch accepts it, the switch enters that information into its topology tables and all connections destined for that NSAP address are routed to that host. These registration messages are sent over the reserved channel VPI 0, VCI 16. You can display the available subcommand by typing `ilmi ?` at the `nsap` level.

```
myswitch::configuration nsap> ilmi ?
  show
```

15.2.1 Displaying NSAP Addresses Registered through ILMI

This command lets you display the NSAP addresses of all of the ports on a switch fabric that have been registered via ILMI. Enter the following parameters:

```
myswitch::configuration nsap ilmi> show
 Port  VPI  NsapAddress                                OrganizationalScope
 1C1    0   47000580ffe1000000f21a2cf00020480697b900   global
 1C2    0   47000580ffe1000000f21a2cf00020480697ba00   global
```

The fields in this display are defined as follows:

Field	Description
Port	The port number on which an NSAP address has been registered via ILMI.
VPI	The virtual path number on which an NSAP address has been registered via ILMI.
NSAPAddress	The NSAP address that has been registered through ILMI for this port.
OrganizationalScope	Shows the availability of this NSAP address to other devices in terms of PNNI routing. The information for this address cannot be distributed outside the indicated scope. Shows the scope for this NSAP address if one has been specified. If no scope has been specified, shows `localNetwork` if the registered NSAP address is a group address and shows `global` if it is an individual address.

To display the NSAP addresses that have been registered via ILMI for a specific port, enter the following parameters:

```
myswitch::configuration nsap ilmi> show 1C1
 Port  VPI  NsapAddress                                OrganizationalScope
 1C1    0   47000580ffe1000000f21a2cf00020480697b900   global
```

If no addresses have been registered via ILMI, the following is displayed:

```
No NSAP address information is available
```

15.3   NSAP-to-E.164 Configuration Commands

These commands let you display NSAP-to-E.164 address mapping information, create an
NSAP-to-E.164 address mapping, and delete an NSAP-to-E.164 address mapping. You can
display the list of available subcommands by typing **e164 ?** at the **nsap** level.

```
myswitch::configuration nsap> e164 ?
  show               new               delete
```

15.3.1 Displaying NSAP-to-E.164 Address Mapping Information

This command enables you to display the current NSAP-to-E.164 address mapping informa-
tion. Enter the following parameters:

```
myswitch::configuration nsap e164> show
 Port VPI NSAP-Address                            Mask Native-E164Address
 3B2  0   47000580ffe1000000f21a00d00020481a00d000 152  4126352756
```

The fields in this display are defined as follows:

Field	Description
Port	The port number on which the NSAP-to-E.164 address mapping exists
VPI	The virtual path number on which the NSAP-to-E.164 address mapping is to be created.
NSAP-Address	The NSAP address for this mapping.
Mask	The number of leading significant bits for this NSAP address.
Native-E164Address	The E.164 address, which can be up to 15 ASCII digits (0-9) long.

If no NSAP-to-E.164 mapping information has been configured, you receive the following
message:

```
myswitch::configuration nsap> e164 show
No NSAP <-> E.164 Mapping information is available.
```

15.3.2 Adding an NSAP-to-E.164 Address Mapping

This command lets you map an NSAP address to an E.164 address format. Enter the following parameters:

```
myswitch::configuration nsap e164> new <port> <vpi> <NSAP> <mask> <E.164>
```

The following is an example of how to add an NSAP-to-E.164 mapping to the mapping table:

```
myswitch::configuration nsap e164> new 3b2 0
0x47.0005.80.ffe100.0000.f21a.00d0.0020481a00d0.0b 152 4126352756
```

15.3.3 Deleting an NSAP-to-E.164 Address Mapping

This command lets you remove an existing NSAP-to-E.164 address mapping. Enter the following parameters:

```
myswitch::configuration nsap e164> delete <port> <vpi> <NSAP> <mask>
```

The following is an example of how to delete an NSAP-to-E.164 mapping from the mapping table:

```
myswitch::configuration nsap e164> del 3B2 0
47000580ffe1000000f21a00d00020481a00d000 152
```

The parameters for new and delete are defined as follows:

Parameter	Description
port	The port number for this NSAP-to-E.164 address mapping.
vpi	The virtual path number for this NSAP-to-E.164 address mapping.
NSAP	The NSAP address for this entry.
mask	The number of leading significant bits for this NSAP address.
E.164	The E.164 address, which can be up to 15 ASCII characters long.

15.4   ILMI LECS Address Registry Commands

These commands let you configure a table of LECS addresses per port. The LAN Emulation Client (LEC) on the user side of the ATM UNI port uses ILMI to query the attached switch for the registered LECS address. You can display the list of available subcommands by typing **?** at the **registry** level.

```
myswitch::configuration nsap registry> ?
  show              delete            new
```

15.4.1 Displaying the ILMI LECS Address Registry

This command lets you display the contents of the ILMI LECS address registry. Enter the following parameters:

```
myswitch::configuration nsap registry> show
Port Type Instance Nsap-Address
1A1  LECS   1      0x47.0005.80.ffe100.0000.f21a.23c0.0020481a23c0.55
1A1  LECS   2      0x47.0005.80.ffe100.0000.f21a.00d0.0020481b7821.00
1A2  LECS   1      0x47.0005.80.ffe100.0000.f21a.116f.002048103086.02
```

The fields in this display are defined as follows:

Field	Description
Port	The port number for the interface. An * to the left of the port number means that entry is still being created and is not active yet.
Type	This is a read-only field that indicates the type of service that is available at the given NSAP address.
Instance	A user-configurable index number that distinguishes between addresses when multiple LECS addresses are assigned to the same port.
Nsap-Address	The LECS address assigned to this port.

If no entries have been configured in the registry, then the following is displayed:

```
myswitch::configuration nsap registry> show
No information is available
```

15.4.2 Deleting an ILMI LECS Address Registry Entry

This command lets you delete an entry or multiple entries from the ILMI LECS address registry. Enter the following parameters:

```
myswitch::configuration nsap registry> delete lecs <port> <instance>
```

For example, you can delete a specific instance on a specific port as follows:

```
myswitch::configuration nsap registry> delete lecs 1A2 2
```

This example only deletes instance number 2 from port 1A2.

You can also delete a specific instance on the control port (CTL) as follows:

```
myswitch::configuration nsap registry> delete lecs 1CTL 1
47000580ffe1000000f21a00d00020481a00d00b
```

This example only delete instance number 1 from port 1CTL. (On an ASX-1000 or a TNX-1100, you need to specify the board number; e.g., 2CTL.)

You can also use an asterisk (*) as a wildcard character to delete all registered addresses on a switch fabric that have the same instance number as follows:

```
myswitch::configuration nsap registry> delete lecs * 2
Are you sure you want to delete *ALL* LECS entries with 2 as the instance [n]? y
```

If you type **n** or press **<ENTER>**, the command is aborted. If you type **y**, all registered addresses on that switch fabric with the instance number 2 are deleted.

Additionally, you can use an asterisk (*) as a wildcard character to delete registered addresses from all ports on a specific network module. The wildcard character does not apply to the CTL port. Enter the * in the network module portion of the port number as follows:

```
myswitch::configuration nsap registry> delete lecs C* 1
```

This example deletes all registered addresses on all ports on network module C with instance number 1.

15.4.3 Adding an ILMI LECS Address Registry Entry

This command lets you add an entry or multiple entries to the ILMI LECS address registry. A maximum of 4 entries may be configured on each port. Enter the following parameters:

```
myswitch::configuration nsap registry> new lecs <port> <instance> <nsap-address>
```

NOTE If the switch port you specify is already configured with an ATM address with the same instance that you specify, the command fails.

For example, you can create a specific instance on a specific port as follows:

```
myswitch::configuration nsap registry> new lecs 1A1 1
47000580ffe1000000f21a00d00020481a00d00b
```

This example only creates instance number 1 on port 1A1.

You can also create a specific instance on the control port (CTL) as follows:

```
myswitch::configuration nsap registry> new lecs 1CTL 1
47000580ffe1000000f21a00d00020481a00d00b
```

This example only creates instance number 1 on port 1CTL. (On an ASX-1000 or a TNX-1100, you need to specify the board number; e.g., 2CTL.)

You can also use an asterisk (*) as a wildcard character for the port number to register the same LECS address on all existing ports on a switch fabric as follows:

```
myswitch::configuration nsap registry> new lecs * 1
47000580ffe1000000f21a00d00020481a00d00b
```

This example creates instance number 1 with the specified address on all existing ports in switch fabric 1.

Additionally, you can use an asterisk (*) as a wildcard character to register the same LECS address on all ports on a specific network module. The wildcard character does not apply to the CTL port. Enter the * in the network module portion of the port number as follows:

```
myswitch::configuration nsap registry> new lecs B* 1
47000580ffe1000000f21a00d00020481b78210a
```

This example registers the specified address as instance 1 on all ports on network module B.

The parameters for delete and new are defined as follows:

Parameter	Description
port	The port number for this ILMI LECS address registry entry.
instance	A user-configurable index number that distinguishes between the LECS addresses when multiple LECS addresses are assigned to the same port. This number can be from 1 to 4.
nsap-address	The NSAP address of the LECS for this entry.

Index